

AD-A228 931

DTIC FILE COPY

4

An Episodic Knowledge Representation for Narrative Texts

Lenhart K. Schubert and Chung Hee Hwang

Technical Report 345
May 1990

DTIC
ELECTE
NOV 23 1990
S B D

DISTRIBUTION STATEMENT A
Approved for public release;
Distribution Unlimited

UNIVERSITY OF
ROCHESTER
COMPUTER SCIENCE

**An Episodic Knowledge Representation
for Narrative Texts**

Lenhart K. Schubert

Chung Hee Hwang

Department of Computer Science
University of Rochester
Rochester, New York 14627, U. S. A.

Department of Computing Science
University of Alberta
Edmonton, Alberta, Canada T6G 2H1

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER 345	2. GOVT ACCESSION NO.	3. RECIPIENT'S CATALOG NUMBER
4. TITLE (and Subtitle) An Episodic Knowledge Representation for Narrative Texts		5. TYPE OF REPORT & PERIOD COVERED technical report
		6. PERFORMING ORG. REPORT NUMBER
7. AUTHOR(s) Lenhart K. Schubert and Chung Hee Hwang		8. CONTRACT OR GRANT NUMBER(s) N00014-82-K-0193
9. PERFORMING ORGANIZATION NAME AND ADDRESS Computer Science Dept., 734 Computer Studies Bldg. University of Rochester, Rochester, NY 14627		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS
11. CONTROLLING OFFICE NAME AND ADDRESS Defense Advanced Research Projects Agency 1400 Wilson Blvd., Arlington, VA 22209		12. REPORT DATE May 1990
		13. NUMBER OF PAGES 53
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office) Office of Naval Research, Inf. Systems Arlington, VA 22217		15. SECURITY CLASS. (of this report) unclassified
		15a. DECLASSIFICATION/DOWNGRADING SCHEDULE
16. DISTRIBUTION STATEMENT (of this Report) Distribution of this document is unlimited.		
17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)		
18. SUPPLEMENTARY NOTES None.		
19. KEY WORDS (Continue on reverse side if necessary and identify by block number) knowledge representation; natural language understanding; episodic logic; situation semantics		
20. ABSTRACT (Continue on reverse side if necessary and identify by block number) (see reverse)		

20. ABSTRACT

We would like to build story understanding systems which are transparent, modular, and extensible. To this end, we have been working on a new logical approach to narrative understanding that features a *GPSG-style grammar* and an *episodic logic* with *probabilistic* inference rules. The grammar represents phrase structure and the relationship between phrase structure and logical form in a modular, explicit form. The logical representation allows propositional attitudes, unreliable generalizations, and other non-standard constructs, providing a uniform, transparent knowledge representation for both the explicit content of stories and for the background knowledge needed to understand them. It makes systematic use of *episodic* variables in the representation of episodic sentences, using these to capture temporal and causal relationships. The rules of inference include probabilistic versions of deduction rules resembling forward and backward chaining rules in expert systems. These can be used for predictive, explanatory, and simulative inference. We illustrate our approach with nontrivial grammar fragments (including semantic rules), and with an extended example of forward-chaining inference based on a sentence from *Little Red Riding Hood*. A pilot implementation is able to make many (though not all) of the inferences we describe.

Accession For	
NTIS GRA&I	<input checked="" type="checkbox"/>
DTIC TAB	<input type="checkbox"/>
Unannounced	<input type="checkbox"/>
Justification	
By	
Distribution/	
Availability Codes	
Dist	Avail and/or Special
A-1	



An Episodic Knowledge Representation for Narrative Texts

Lenhart K. Schubert
Chung Hee Hwang

Abstract

We would like to build story understanding systems which are transparent, modular, and extensible. To this end, we have been working on a new logical approach to narrative understanding that features a *GPSG-style grammar* and an *episodic logic* with *probabilistic* inference rules. The grammar represents phrase structure and the relationship between phrase structure and logical form in a modular, explicit form. The logical representation allows propositional attitudes, unreliable generalizations, and other non-standard constructs, providing a uniform, transparent knowledge representation for both the explicit content of stories and for the background knowledge needed to understand them. It makes systematic use of *episodic* variables in the representation of episodic sentences, using these to capture temporal and causal relationships. The rules of inference include probabilistic versions of deduction rules resembling forward and backward chaining rules in expert systems. These can be used for predictive, explanatory, and simulative inference. We illustrate our approach with nontrivial grammar fragments (including semantic rules), and with an extended example of forward-chaining inference based on a sentence from *Little Red Riding Hood*. A pilot implementation is able to make many (though not all) of the inferences we describe.

(21)

A brief preliminary version of this paper was presented at *KR'89* (Schubert and Hwang 1989).

1 Introduction

Many ideas and systems have been developed for narrative understanding, and some of them, e.g., CYRUS (Kolodner 1981) and BORIS (Lehnert *et al.* 1983, Dyer 1983), have shown a remarkable degree of understanding in complicated human domains. These ambitious systems, however, are very complex and hard to extend beyond the few stories they handle. We believe this is so because of insufficient *expressiveness* and *clarity* of the knowledge representations used, for knowledge about language as well as about the world, and about inferences warranted by that knowledge.

For example, the meaning representations used often cannot express complex quantification ("most people with two or more cars"), logical compounding ("If he fails, he is either lazy or a fool"), complex concepts ("the type of person who never forgets a slight"), modification ("a nearly invisible pale brown birthmark"), temporal relations ("He had seen her twice the previous week"), and so on. As well, knowledge about language and about the world are often buried in procedures (e.g., procedures which seek semantically appropriate fillers for frame slots) in a way that makes it very hard to determine what linguistic and factual assumptions have been made.

These considerations have led us to an approach to narrative understanding in which all types of linguistic, world and inference knowledge are represented in an explicit, analyzable form. We have chosen Generalized Phrase Structure Grammar (GPSG; Gazdar *et al.* 1985) as our grammatical representation and have been developing a knowledge representation, called *episodic logic* (a kind of situational logic, *cf.*, Barwise and Perry 1983, Barwise 1989), for encoding both the content of narratives and the knowledge needed to understand them. GPSG is a particularly perspicuous grammatical formalism which is expressively adequate for almost all English grammatical phenomena, and is relatively easy to use by a parser and logical-form generator. Episodic logic is expressively rich and close enough to surface form so that the relationship between surface form and logical form can be specified in a modular, transparent way. It introduces episodic variables so that implicit, context-dependent relationships among episodes (events, situations, circumstances, eventualities, *etc.*) can be made explicit. It also allows the representation of restricted quantifiers, propositional attitudes, predicate modifiers, nominalized predicates, and perhaps most importantly, unreliable generalizations. Such generalizations have recently received much attention in the non-monotonic reasoning literature and elsewhere (e.g., linguistic semantics). The practical adequacy of our logic has been tested on small story fragments.

In the next section, we motivate and explain some of the unusual features of our logic. In section 3, we provide preliminary semantics. While our main concern in this paper is with representation and inference, we need to make plausible the claim that linguistic input could be mapped into our logical representation in a principled and transparent fashion. Therefore, in section 4, we sketch the derivation of episodic logical form from surface structure using a GPSG grammar. Next, in section 5, we introduce some inference rules and indicate their role within our implementation. In section 6, we

illustrate the inference process in story understanding with an extended example based on a small fragment of *Little Red Riding Hood*. In the concluding section, we comment on related work and assess the progress made and work still to be done.

2 Episodic Logic

In our approach to interpreting English text, the representation of an input sentence is obtained in several (possibly interleaved) processing stages. Initial representations are in general ambiguous (e.g., with respect to the scopes of quantifiers and other operators) and context-dependent (e.g., involving indexical operators like *past*, whose interpretation depends on the utterance time). The various processing stages are aimed at removing ambiguity and context-dependence. We will loosely refer to the outputs of any of the stages as “logical forms.” However, the emphasis in this section is on the *final* logical forms, i.e., the context-independent representations which are ultimately “committed to memory,” ready for use in inferential processes.

Our initial logical form closely follows surface structure, and even the final form is rather English-like. It “mimics” noun phrases in its use of restricted quantifiers, and follows English sentence syntax by having the “subject” of a predication precede the predicate. For example, “Every dog has a tail” would be represented (minus tense) as

- (1) $(\forall x:[x \text{ dog}] (\exists y:[y \text{ tail}] [x \text{ have-as-part } y]))$

after predicate disambiguation and quantifier scoping.¹ (The initial representation computed from the surface form would be $[<\forall \text{ dog}> \text{ have } <\exists \text{ tail}>]$.)

Two more features that lead to close conformity between the surface form and our logical form are predicate modification and λ -abstraction. These are illustrated by (2) and (3), with tense again neglected (and with the phrase in (3a) represented in unscoped form in (3b)):

- (2) a. Canada is very distant from Australia
 b. $[\text{Canada (very } \lambda x[x \text{ distant-from Australia}])] ;$
 (3) a. the brother of Mary who is a doctor
 b. $<\text{The } \lambda x[[x \text{ brother-of Mary}] \wedge [x \text{ doctor}]]>.$

In (2b), the predicate modifier *very* is a function which, when applied to a predicate, yields another, more restricted predicate. Other examples of predicate modifiers are *almost*, *fake*, *coll* (forming predicates over collections), and complex modifiers corre-

¹We use restricted quantification of the form $(Q\alpha:\Phi\Psi)$, where Q is a quantifier, α is a variable, and Φ and Ψ are formulas. That is, $(\forall\alpha:\Phi\Psi)$ and $(\exists\alpha:\Phi\Psi)$ are equivalent to $(\forall\alpha)\Phi \rightarrow \Psi$ and $(\exists\alpha)\Phi \wedge \Psi$, respectively. When there is no restriction Φ , we write $(Q\alpha \Psi)$. Also note that we use square brackets to indicate predicate infix expressions, round brackets for prefix expressions, and angle brackets for unscoped operators. Scoping of quantifiers and other operators is discussed in Schubert & Pelletier (1982), Hurum & Schubert (1986), and Hurum (1987,1988).

sponding to adverbials such as (*adv-a* $\lambda a[a$ with-instrument key3])² which means “with the key.”

Several operators for nominalizing (reifying) sentence or predicate intensions are also available in episodic logic, including a proposition-forming operator *That* and kind-forming operators (*K*, *K*₁ and two others) as illustrated in (4) to (7) (again neglecting tense)

- (4) a. That Mary smoked is not true
b. $\neg[(\text{That } [\text{Mary smoke}]) \text{ true}]$;
- (5) a. Mary knows that John is intelligent
b. $[\text{Mary know } (\text{That } [\text{John intelligent}])]$;
- (6) a. Snow is white
b. $[(K \text{ snow}) \text{ white}]$;
- (7) a. A dog is a mammal
b. $[(K_1 \text{ dog}) \text{ mammal}]$.³

Another distinctive feature of our logic, responsible for its name, is the inclusion of *episodic variables*. Whereas examples (1) - (2) dealt with “enduring” properties (though not *eternal* ones, a fact for which we eventually allow), (8) and (9) below involve “episodic” ones:

- (8) Everyone looked at Mary;
- (9) This (event) made her blush.

The final representations of (8) and (9) involve a connective “**” relating a sentence to the episode it *characterizes* (or *completely describes*). With some simplifications, they are:

- (8') $(\exists e1:[e1 \text{ before } \text{Now1}]$
 $[(\forall x:[x \text{ person}][x \text{ look-at Mary}]) ** e1])$;
- (9') $(\exists e2:[e2 \text{ before } \text{Now2}]$
 $[[e1 \text{ cause-of } e2] \wedge [[\text{Mary blush}] ** e2])$.

Thus (8') says that *e1* is an episode characterized (or completely described) by “Everyone looked at Mary,” and similarly for *e2* in (9').⁴

²*adv-a* (standing for *a*ction-modifying *a*dverbial) is an operator that forms a predicate modifier from a predicate over actions such as *passionate* or $\lambda a[a$ for-benefit Mary] which means “for Mary.”

³*K* is used in the interpretation of mass nominals like *snow* and non-numeral bare plurals like *dogs*, and *K*₁ in the interpretation of indefinite count singulars like *a dog* and bare numeral plurals like *twelve eggs* (cf., μ and μ_1 in Schubert and Pelletier 1987).

⁴Note the reduction of past tense to a relation placing episodes *e1* and *e2* before *Now1* and *Now2* (more accurately, before the utterances of the sentences (8) and (9)), respectively. This reduction is obtained from an initial translation involving indexical operator *past* – see section 4. Also, (8') would be further expanded to show *individual* episodes of “person *x* looking at Mary,” occurring *during* the

A characterizing description of an episode provides *all* the facts that are supported by it, except possibly for ones entailed by those given, via meaning postulates (MPs). For example, blushing entails becoming (more) red in the face, and so “Mary blushed, becoming red in the face” (represented logically) could still characterize episode e_2 . By contrast, “Mary blushed, feeling foolish” cannot characterize the same episode, only a factually more comprehensive one, since it says strictly more (i.e., feeling foolish is not entailed by blushing, via MPs). On the other hand, “Mary’s face changed color” cannot characterize the same episode either, only a factually less comprehensive one, since it says strictly less (i.e., one’s face can change color, e.g., by turning pale, without one’s blushing). As we will shortly indicate, the notion of a characterization (or complete description) is important for getting *causal* statements right.

Our logic also contains a weaker but more fundamental operator “*” that reads “partially describes.” “*” is essentially an object-language embedding of the semantic notion of *truth over an episode or situation*. (As such, $[\Phi * e]$ is similar to $[\Phi]^*(e)$ in (Reichenbach 1947); and to $e \models \Phi$ in (Barwise 1989).) For example, $[[\text{Mary blush}] * e]$ entails the truth of [Mary blush] in episode e . As in the case of $[[\text{Mary blush}] ** e]$, the blushing must extend over the entire episode. The following meaning postulate relates “**” to “*” (where Φ and Ψ are schema variables over formulas, and η is a schema variable over terms):

$$\Box [[\Phi ** \eta] \rightarrow [[\Psi * \eta] \rightarrow \Box (\forall e [[\Phi * e] \rightarrow [\Psi * e]])]],$$

i.e., if Φ characterizes η , then if Ψ partially describes η , it does so as a matter of necessary consequence from Φ .

To see the significance of the distinction between “characterizations” and arbitrary partial descriptions of episodes, suppose that the “**” in (8') were replaced by “*.” Then (as a little thought shows) (8') and (9') would be true in a situation in which everyone looked at Mary, *laughing derisively*, and it was *this* (more complex) event that made her blush. Yet we would not say that (8) and (9) are both true in such a situation; so (8'), with “**” weakened to “*,” would not be a correct formalization of (8). To put it schematically, the inference

$$[\Phi ** e_1], [\Psi ** e_2], [e_1 \text{ cause-of } e_2] \vdash [[\Psi ** e_2] \text{ because } [\Phi ** e_1]]$$

is sound, while

$$[\Phi * e_1], [\Psi * e_2], [e_1 \text{ cause-of } e_2] \vdash [[\Psi * e_2] \text{ because } [\Phi * e_1]]$$

is not.

Another point to be noted above is the *free* occurrence of variable e_1 in (9'), outside the scope of its quantifier. This is permissible in our logic in a sequence of conjoined sentences or in a conditional, thanks to a “parameter” mechanism similar to

overall episode e_1 . Finally, the representation of (9) would actually contain the logical translation of the verb phrase *make Mary blush*, viz., $((\text{make blush}) \text{ Mary})$, from which (9') would be derived by meaning postulates.

that of discourse representation theory (DRT) (cf., Kamp 1981, Heim 1982). In effect, the conjunction of (8') and (9') would be interpreted as if the \exists -quantifiers had widest scope (see the semantics of " \exists ," "The," " \wedge " and "&" in the next section).

The parameter mechanism is also the key to representing generic conditionals, such as "A wolf is (usually) gray," "A child (usually) loves his or her grandmother," or "When two strangers meet in a deserted region, they often greet." Generic conditionals take the following form:

$$(\exists x_1(\exists x_2(\dots(\exists x_k \Phi) \dots))) \rightarrow_p \Psi,$$

where Φ and Ψ are formulas involving x_1, x_2, \dots, x_k , and p is a numeric lower bound on frequency (statistical probability). The semantics of the connective " \rightarrow_p " essentially allows it to "take control" of parameters (existentially quantified variables) in the antecedent and iterate over their denotations. Thus,

$$(\exists x[x \text{ wolf}]) \rightarrow_{.8} [x \text{ gray}]$$

says that a wolf is usually gray (or, most wolves are gray). Generic conditionals are often used in *causal* axioms. In particular, *predictive* causal axioms assume the occurrence of some particular type of episode $e1$ in the antecedent, and predict another episode $e2$ caused by $e1$ in the consequent. The following is an example of a predictive axiom:

*When a predatory animal sees a non-predatory creature of comparable or smaller size, it may want to attack and eat it.*⁵

$$\begin{aligned} & [(\exists x:[x \text{ predatory-animal}] \\ & \quad (\exists y:[y \text{ creature}] \wedge (\neg[y \text{ predatory-animal}]) \wedge \\ & \quad \quad [[y \text{ as-big-as } x] \vee [y \text{ smaller-than } x]]) \\ & \quad (\exists e1[[x \text{ see } y] ** e1])))] \\ & \rightarrow_{.6} [(\exists e2:[[(\text{begin } e2) \text{ during } e1] \wedge [e1 \text{ cause-of } e2]] \\ & \quad \quad [[x \text{ want } (K_a (\text{attack } y))] ** e2]) \wedge \\ & \quad (\exists e3:[[(\text{begin } e3) \text{ during } e1] \wedge [e1 \text{ cause-of } e3]] \\ & \quad \quad [[x \text{ want } (K_a (\text{eat } y))] ** e3]])] \end{aligned}$$

Equally important are *explanatory* axioms, such as the following:

If a creature wants to eat some food, it is likely to be hungry.

$$\begin{aligned} & [(\exists x(\exists e1[[x \text{ want } (K_a \lambda y(\exists z:[z \text{ food-for } y] [y \text{ eat } z]))] ** e1))] \\ & \rightarrow_{.9} (\exists e2:[e2 \text{ cause-of } e1] \wedge [e2 \text{ same-time } e1]) \\ & \quad [[x \text{ hungry}] * e2)] \end{aligned}$$

⁵ K_a below is an operator that forms a kind of action from an action predicate intension. More on this later.

Much of the world knowledge we use in our experimentation is in fact stated as causal axioms like these.

We indicated at the outset that episodes, as we construe that term, are the same as events, situations, circumstances or eventualities (though we may eventually want to treat "events" as being episodes in which something *happens*, as opposed to those in which some state or process persists). We carefully avoided including *actions* or *activities* in this list since we do not regard these as being of the same type as episodes. For example, it seems to us that (10) and (11) may describe the same episode or event (an exchange of a car for a sum of money), but different actions (a buying and a selling)

(10) John bought the car from Mary;

(11) Mary sold the car to John.

Note, in particular, that the buying in (10) may have been performed *reluctantly* and the selling in (11) *eagerly*, but it would be very odd to say that the *events* described in (10) or (11) were reluctant, or eager, or occurred reluctantly or eagerly. Events simply do not have such properties.

Nevertheless, events and actions seem very closely related. In particular, specific actions appear to determine specific events (namely, the *performances* of those actions). The difference between them, we suggest, is just that actions have well-defined agents while events do not. Thus it makes perfect sense to ask about an action "Who did it?" (i.e., "Who performed the action?"), while it is odd to ask about an event "Who did it?" (i.e., "Who performed the event?")

If specific actions determine specific events, and actions have well-defined agents while the events they determine do not, then the simplest possible theory of their relationship is this: actions are ordered agent-event pairs. Thus, for an event e which is a performance of an action a by agent x , $a = [x \mid e]$, where " \mid " is the pairing function. The agent of an action is then just the first element of the action, written ($\text{fst } a$), and the event of the agent performing the action is the second element of the action, written ($\text{rst } a$). We have found this view very helpful in the analysis of adverbials (some of which we regard as modifying events – e.g., locative and temporal adverbials – and some of which we regard as modifying actions – e.g., manner adverbials). Also, the distinction has resolved some persistent difficulties we encountered in reasoning about actions, such as the "wicked" actions of the wolf in the story of *Little Red Riding Hood*.

We should remark that while actions, on our account, are agent-event pairs, the converse is not in general true: many agent-event pairs are not actions. For instance, if e is the event of the sun rising, then $[\text{John} \mid e]$ is certainly not an action. Whether or not an agent-event pair is an action depends entirely on the characterization of the event, a dependence that is to be captured by meaning postulates.

In this connection, we should mention the two remaining kind-forming operators, K_e and K_a , forming kinds of events and kinds of actions (more generally, *attributes*) from sentence intensions and (action) predicate intensions respectively as shown in the

following examples:

- (12) a. For Mary to dance was rare
 b. $[(K_e [\text{Mary dance}]) \text{ rare}]$;
- (13) a. To kiss Mary is fun
 b. $[(K_a (\text{kiss Mary})) \text{ fun}]$.

Finally, some preliminary remarks are in order about the inference rules we use in narrative understanding. One very general inference rule, resembling those used in expert systems, is called **Rule Instantiation (RI)** (see section 5 for details, where the dual of RI, called **Goal Chaining**, is discussed as well). In some cases, this can be thought of as a general form of *modus ponens* with universal instantiation and use of multiple minor premises (instantiating the antecedent of a universally quantified conditional). However, it also allows instantiation of generic conditionals such as the predictive and explanatory axioms above. As a simple example, RI allows the inference

$$\frac{(\exists x[x \text{ wolf}]) \rightarrow_s [x \text{ gray}], \quad [W \text{ wolf}]^9}{[W \text{ gray}]^{.72}}$$

The superscripted numbers are interpreted as lower bounds on *epistemic* probabilities (in contrast with the *statistical* interpretation of probabilities modifying the connective in generic conditionals).

Note that successive inference steps of this type will lead to attenuation of the probabilities assigned to successive conclusions. Such inference chaining (with safeguards against certain fallacies) is done routinely in our implementation. A more problematic issue is the "parallel" combination of evidence. This issue arises when several generalizations (or inference chains) assign different epistemic probabilities to the same formula. A partial solution, applicable when the antecedent of one generalization is more specific than the other, i.e., entails the other, is to apply only the more specific rule (*cf.*, Kyburg 1983, Bacchus 1988). However, this leaves open the question of how to combine logically independent (or only "probabilistically dependent") bits of evidence. We hope to apply methods similar to those of Pearl (1988), though these are aimed essentially at sentential reasoning and so will require considerable extension.

3 Semantics

In an attempt to develop more adequate meaning representations (and knowledge representations) for NLU systems, there is a trade-off between scope and rigor. One can either cover many of the semantic phenomena of language in a rough-and-ready fashion, or a few of them with a completely worked out formal framework. (For some sense of the work left to be done in situation semantics – even without detailed consideration of linguistic matters – see Barwise's (1989) essay on "branch points.") As AI researchers

interested in building working NLU systems, we believe it would be a mistake to cut back on expressive power for the sake of theoretical tractability. Even the simplest stories confront us with all the traditional conundrums – tense, aspect, adverbials, generic sentences, and so on. The best strategy, for our purposes, is a “top-down” one, using a representation with more or less full coverage of English from the outset, and to subject this representation to continual revision in the light of the interlocking needs of grammar, computation of logical form, formal semantics, and inference.

Inevitably, then, many uncertainties and gaps remain in our logical semantics. However, what is worthwhile about our attempt is that it gives detailed expression to the intuitions which are guiding our work, subsumes classical logic, provides tentative extensions in several major directions, and is sufficiently carefully formalized to make future systematic analysis and revision possible. Moreover, the semantic ideas were not conceived in isolation, but with an eye on the mapping from surface structure to logical form, and on the inferences that support story understanding. In these respects, the attempt is at least a step in the right direction.

Episodic logic is in some ways a sequel to earlier efforts by Schubert and Pelletier (1982, 1989) to provide a first-order logical form for English, including generic sentences. As before, the attempt is Montague-inspired, avoids higher-order predicates (*cf.*, Chierchia and Turner 1988), and incorporates a DRT-like treatment of indefinites. However, the new logic relies crucially on the episodic operators and variables introduced in the last section and grounds intensions in possible situations (or episodes) rather than possible worlds. Also, it moves away from the amalgamation of contexts and interpretations in Schubert and Pelletier (1989) (more on this later).

Our semantics of situations has naturally been influenced by situation semantics (Barwise and Perry 1983; Barwise 1989), even though we do not share situation semanticists’ distaste for possible situations (or other possible entities), logical form, and meaning postulates. From our rather practical point of view, their scruples seem a little perverse. After all, situation semantics embraces “nonactual situations” supporting nonactual states of affairs/facts/infons; employs a language of relational tuples, type abstraction, support (of facts by situations) and constraint relations with striking resemblances to familiar logical forms; and posits constraints which are similar to meaning postulates in all but name. We think that possible situations provide a simple basis for the semantics of propositional attitudes, that logical form is indispensable in building NLU systems, and that “meaning postulates” is a good term for those general statements which are taken to be true in virtue of the meaning of their atomic constituents.

In any case, the essential idea we have adopted from situation semantics is that sentences are used to describe situations (episodes, events, . . .), which are limited pieces of reality in terms of the objects they involve, the time they span, and the sentences whose truth they support. We implement this idea by treating semantic interpretations, and the valuation functions that extend them, as *partial functions on situations* (*cf.*, Fenstad *et al.* 1987). In contrast with situation semantics, episodic logic does not treat

properties as primitive, though it does admit (reified) properties as individuals. Also, we see no point in trying to treat "the facts of a situation" uniformly as relational tuples. For instance, the facts $\{1 \neq 2, 2 \neq 1\}$ seem to be the same, yet the tuples $\langle =, 2, 1, no \rangle$, $\langle =, 1, 2, no \rangle$ are distinct. Similarly, $\{1=1 \vee 1=2, 1=2 \vee 1=1, 1 \text{ has the property of being identical with } 1 \text{ or } 2\}$ all seem to be the same fact, yet $\langle \vee, \langle =, 1, 1, yes \rangle, \langle =, 1, 2, yes \rangle, yes \rangle$, $\langle \vee, \langle =, 1, 2, yes \rangle, \langle =, 1, 1, yes \rangle, yes \rangle$, and $\langle [n \mid \langle \vee, \langle =, n, 1, yes \rangle, \langle =, n, 2, yes \rangle, yes \rangle], 1, yes \rangle$ are all distinct. (That is not to say we have a full account of facts, but we think they should be in one-to-one correspondence with true propositions, where propositions in turn are in one-to-one correspondence with sentence intensions.)

To speak of *sentence* intensions, and even of *sentences* describing situations (or situation types), is not necessarily loose talk. Consider a context-dependent sentence such as "I am with you." It is usual in possible worlds semantics to say that a specific intension for such a sentence can be reached only by way of an *index* (context) which supplies the time and place of speech, a speaker, addressee, and other things. However, it is perfectly possible to put the burden of supplying all this information squarely on the interpretation function, i.e., the function that supplies basic intensions. So this function would supply well-defined denotations not only for "Socrates" and "Picasso" but also for "I," "you," and "here" and "now," and well-defined intensions not only for "love" and "walk," but also for "be with" (which can mean "be near," "side with," and other things), and so on. This is a well-known option, and was used in Schubert and Pelletier (1989), where interpretation functions, extended to valuation functions, were called *contexts*. (The same idea also seems to be implicit in Barwise's (1989) "perspectival situations.")

However, this route is not practical for computational purposes. Suppose we were to represent an input sentence like the above essentially as received, e.g., as (pres [I be-with you]). This would be quite useless for inference in conjunction with other formulas obtained in other circumstances (so that *pres*, *I*, etc., would have no *fixed* meaning). At the very least, we would need to annotate these indexical formulas with information about the contexts in which they were obtained, so that the *pres* and *I* (etc.) of one context would not be confused with the *pres* and *I* of another context, featuring a different speech time and speaker.

But even such annotated formulas would be impractical for inference, as far as we can tell. The amount of contextual information that would have to be included in the annotation would be quite large (including substantial amounts of prior discourse to provide unambiguous referents for anaphora, description of external objects and circumstances to provide referents for deictic phrases, etc.), and the computational burden would be correspondingly heavy. Hence we have chosen to follow the usual practice in the design of NLU systems of assuming that the initial, indexical representations of inputs are converted to a nonindexical ones, using the contextual information which would otherwise have to be permanently attached to their representation. Note that "deindexicalized" representations of inputs can again be said to describe situations, and to have intensions,




Figure 1: Ontology of basic individuals

but this time they do this in accord with a *fixed* interpretation function, rather than one that changes from moment to moment. Thus deindexalized representations will be usable in concert for inference, without regard for their origins. The following discussion refers to this non-indexical form of episodic logic.

Model structures for episodic logic are based on a very liberal ontology of *possible individuals* \mathcal{D} . Our aim is to include in \mathcal{D} everything we can talk about. Thus, besides “ordinary” individuals, we include disjoint sets of *possible situations* \mathcal{S} , *propositions* \mathcal{P} , and *properties* \mathcal{Q} , and *kinds, collections and n -tuples* ($n = 2, 3, \dots$) of all of these. (Kind-, collection- and tuple-formation may be iterated, i.e., there are kinds of kinds, collections of collections, kinds of collections, kinds of kinds of collections, and so on.) The reason for allowing “possible individuals” is that ordinary talk abounds with them, as in “Sherlock Holmes is a fictitious detective” and “Tomorrow’s lecture has been cancelled” (see Hirst 1989). Note that the latter sentence, if true, refers to a nonexistent (but possible) event. Figure 1 shows the assumed relations between the basic categories of individuals.

Total functions *begin*: $\mathcal{S} \rightarrow \mathcal{R}$ and *end*: $\mathcal{S} \rightarrow \mathcal{R}$ determine the beginning and ending “clock times” of all situations, where \mathcal{R} is a set of numbers (possibly the reals or rationals), along with $-\infty$ and $+\infty$. For all situations s , $\text{begin}(s) \leq \text{end}(s)$. We say s is *during* s' iff $\text{begin}(s') \leq \text{begin}(s) \leq \text{end}(s) \leq \text{end}(s')$.

Two disjoint relations over $\mathcal{D} \times \mathcal{S}$, namely, *Actual* and *Nonactual*, determine what entities are actual and nonactual relative to a situation. Together, these are the *participants* in the situation. Situations are partially ordered by a part-of (or subepisode) relation “ \sqsubseteq ,” where $s \sqsubseteq s'$ implies *Actual*(s, s'). Moreover, for each $s \in \mathcal{S}$ there is at least one maximal element $w \in \mathcal{S}$ (called a *possible world*) with $\text{begin}(w) = -\infty$, $\text{end}(w) = +\infty$, such that $s \sqsubseteq w$. The corresponding join operator \sqcup forms a join semilattice over the situations which are *Actual* relative to any given world w . \sqcup is understood as

joining two episodes into a (temporally and/or factually) larger episode. (Note that this can lead to episodes with temporally disjoint subepisodes.) With respect to this operator, clock times obey the rules $begin(s \sqcup s') = \min\{begin(s), begin(s')\}$ and $end(s \sqcup s') = \max\{end(s), end(s')\}$, whenever $s \sqcup s'$ is defined.⁶ Consequently, subepisodes occur during (possibly, concurrently with) the episodes of which they are a part, i.e., for all $s, s' \in \mathcal{S}$, $s \sqsubseteq s'$ implies s during s' . For example, the event of John's eating the strawberries greedily has as a concurrent subepisode his eating the strawberries (i.e., this subepisode lacks some of the specifics of the original episode), and it has John's swallowing the strawberries as a temporally partial subepisode (since as a matter of meaning, eating something involves swallowing it). With "during" defined as above, \sqsubseteq can be specialized to a "concurrent part" (or concurrent subepisode) relation " \preceq ," i.e., $s \preceq s'$ iff $s \sqsubseteq s'$ and s during s' and s' during s .

We can now identify *possible times* \mathcal{I} as those situations (n.b., *situations*, not numeric clock times) maximal in the \preceq -ordering; i.e., they have "maximal content" relative to their *concurrent* parts. Times, in this sense, are arbitrarily short (possibly momentary) or long (possibly unbounded) time *intervals*, such as "this moment," "the year 1990," or "the history of the universe." Their factual content is "everything that happened, or was the case, during that time." Among the possible times, the temporally unbounded ones are thus the *possible worlds* \mathcal{W} ;⁷ the temporally minimal ones are referred to as the *moments of time* \mathcal{T} . Note that \mathcal{W} and \mathcal{T} are the traditional indices of possibility. Two functions that yield moments from times are *init* and *fin*, where *init*(i) is the initial moment of i , and *fin*(i) is the final moment of i . These functions are assumed to be total on \mathcal{I} . Individuals can be *Actual* relative to any number of worlds, and must be *Actual* or *Nonactual* relative to any *given* world. In general, situations, like other individuals, can be *Actual* relative to any number of worlds; however, times are *Actual* relative to exactly one world, namely, the one of which they are a part, and *Nonactual* relative to all others. Accordingly, we assume a function *world*(i) supplying the unique world of a given time i . Intuitively, the reason for this assumption is that times, being factually (propositionally) maximal relative to all concurrent situations, already encapsulate the history of their universe; i.e., at any given time i , the fact that some earlier event occurred is a *fact of i*, and similarly for future events. So on our conception of times, the world component of a world-time index is redundant, and indeed in our semantics there is just one index of possibility, a situational one.

We have also positioned a set of *possible facts* \mathcal{F} in our ontology, namely, the temporally unbounded situations (including the possible worlds, which are thus "maximally inclusive possible facts"). So, the situation that $2 + 2 = 4$, being atemporal (temporally unbounded), counts as a fact. Actual facts are surely related to true propositions (as

⁶The binary join can be extended to a set join, where we assume that for a set of situations \mathcal{E} which are *Actual* relative to some world w , $begin(\sqcup \mathcal{E}) = glb\{begin(s) | s \in \mathcal{E}\} \in \mathcal{R}$, and $end(\sqcup \mathcal{E}) = lub\{end(s) | s \in \mathcal{E}\} \in \mathcal{R}$.

⁷The universe may in a mathematical sense be temporally bounded, but on our intuitive reckoning, it cannot be. The first question that comes to mind about the Big Bang is "What went on *before* that?"

we previously suggested, in one-to-one fashion), but as Vendler (1967) argued, cannot be identified with them. For instance, facts can be discovered and can be due to other facts, but propositions cannot (only the *fact* of their truth or falsity can).⁸

We leave open the exact nature of facts and propositions, but assume a total function $' : \mathcal{S} \rightarrow \mathcal{F}$, giving the facts of a situation s as $'s$. For a fact f , $'f = f$, and for a time i , $'i = 'world(i)$. Also, we assume a one-to-one partial function $\wedge : (\mathcal{S} \rightarrow 2) \rightarrow \mathcal{P}$, (interpreting *That*) for forming propositions from sentence intensions (here $(\mathcal{S} \rightarrow 2)$ is the set of partial functions from \mathcal{S} to 2, 2 being the truth values $\{0, 1\}$); its inverse is $\vee : \mathcal{P} \rightarrow (\mathcal{S} \rightarrow 2)$, (interpreting *True*) for forming a sentence intension from a proposition). Thus, $\vee \wedge \sigma = \sigma$ for sentence intension σ , and $\wedge \vee \pi = \pi$ for proposition π . If $\underline{*}$ is the function interpreting the characterizing relation, so that $\underline{*} : (\mathcal{S} \rightarrow 2) \rightarrow (\mathcal{S} \rightarrow (\mathcal{S} \rightarrow 2))$, it appears that the fact corresponding to proposition π is $(\iota f (\vee \pi \underline{*} f))$, and the proposition corresponding to fact f is $(\iota \pi (\vee \pi \underline{*} f))$.

A basic choice that needs to be made in characterizing an interpretation I and its extension to a valuation function $\llbracket \cdot \rrbracket_I$ is whether types of situations are "outward persistent," i.e., whether situations (episodes) are of the same types as their temporally smaller parts (subepisodes). If an episode consists of John eating his strawberries and then drinking his coffee, should we regard the *entire* episode as being of (tenseless) type "John eat his strawberries" and "John drink his coffee"? I.e., does the episode as a whole support the truth of the (tenseless) subepisode descriptions?

It is tempting to answer affirmatively, since truth-functional compounds are then more easily evaluated. For instance, a conjunction will be true in a time (interval) just in case both conjuncts are true in it (modulo the necessary provisions for DRT-like handling of indefinites). In a nonpersistent approach, we instead have to say that the conjunction is true over a time (interval) just in case each conjunct is true over a subinterval, and the two subintervals together "span" the given interval. As well, seemingly valid arguments like the following suggest outward persistence: John resigned today; therefore, he resigned this week. However, both kinds of evidence are problematic. First, an outward-persistent semantics in which Φ is true of a time interval whenever it is true of a subinterval would obliterate information. Suppose, for instance, that John solved a mathematical problem over a time period i , and in the course of doing so solved another (smaller) problem over subinterval j . Then the persistent semantics obliterates the larger problem-solving episode as a distinguishable instance of "John solved a mathematical problem," since that description already follows by persistence from the smaller instance over interval j . This becomes important when one applies temporal qualification such as "John solved a mathematical problem, and immediately wrote down the solution," or "John solved a mathematical problem. It took him five hours." These qualifications do not have their intended force if the episode to which they

⁸Situations support sets of facts (or the truth of sets of propositions), but since facts can also be facts about situations (and thus indirectly about sets of facts), not all sets of facts can be the facts of some situation (by cardinality).

are applied could be one which was arbitrarily “inflated” (via outward persistence) to a longer episode.

More generally, much as we might prefer a simple boolean algebra as a basis for the truth conditions of all logically compound episode descriptions, truth conditions for tenseless (and therefore potentially temporal)⁹ sentences simply do not appear to be boolean. For instance, $[\Phi \wedge \Phi]$ apparently does not entail Φ (for temporal sentence Φ). Given that “Over three years, Mary had a baby, and a baby, and a baby,” the most natural interpretation is one in which the overall episode described is one containing three separate subepisodes of Mary having a baby. Yet the overall episode is not one of Mary’s having a baby; rather, it is one of Mary having *three* babies. At the same time we should emphasize that on our view, truth values of *atemporal* sentences *do* combine truth-functionally at any given time (or world). Thus nonindexical predications of the sort “Mary has a baby in episode *E17* before *Now12*, and Mary has a baby in episode *E18* before *Now12*, etc.,” will combine in the accustomed boolean fashion.

As far as the above linguistic evidence for outward persistence is concerned, it carries little weight, since *today* should in any case be interpreted as “during today” (which of course entails “during this week”). For if we take “John resigned today” as describing a day-long episode (rather than one *during* today), we will have great trouble with “John resigned today in a five-minute meeting with the boss.”

Thus, we opt for non-outward persistent temporal situation types, and hence for truth conditions considerably more complex (but also, we claim, more accurate) than those usually proposed in situation semantics. Atemporal situation types, however, are outward persistent. Moreover, we will assume a more restricted sort of persistence for *all* situation types, including temporal ones, namely “upward persistence.” By this we mean persistence through the \preceq relation, i.e., from a situation to a *concurrent* one with more “information.” This seems natural to us despite the following sort of objection. Suppose it is true that “Everyone worked hard this morning” if understood as referring to today’s events at the office. Shouldn’t the same sentence become false once we expand our purview to include *more of the world*? Reasonable as this objection is, we think concurrent upward persistence nevertheless holds for *deindexicalized* representations. In such a representation, a phrase like *everyone* would have been replaced by something like *everyone at the office*, obviating any need for repeated inference about the sense of *everyone* that was intended at the time of its use (cf., Kratzer 1987). The assumption of upward persistence allows us to avoid further complications in the specification of interpretations and valuations. Namely, it allows us to use values of expressions at times $i \in \mathcal{I}$ to impose more general constraints on values at situations $s \in \mathcal{S}$, via upward persistence.

⁹The fact that only untensed sentences can be temporal needs to be clearly understood: [Mary marry John], corresponding to tenseless sentence “Mary marry John,” is temporal in the sense that it applies truthfully to some times (at which marrying occurs) but not to others. By contrast, [[Mary marry John] * E], as part of the deindexicalized representation of tensed sentence “Mary married John,” is *atemporal* (a fact or falsehood); it is true at all times (and for all time), or at none.

Before stating what constitutes an interpretation I and valuation function $\llbracket \cdot \rrbracket_I$, we should mention that our semantic rules assume “curried” predicate and function application, in prefix form. The “flattened” infix form of predication of section 2 (with infix expressions in square brackets) is obtained by some rather obvious flattening (i.e., bracket deletion) rules and an infixing rule that places the *last* argument of a predicate in “subject” position, ahead of the predicate. (However – following Reichenbach (1947) – we take the sentential argument Φ in $[\Phi ** \eta]$, $[\Phi * \eta]$, and $[\Phi @ \eta]$ to be first, semantically as well as syntactically. Also, the pairing function “|” and dyadic arithmetic functions may be infix, with the function symbol following its *first* argument.)

Also, we should emphasize that we are not giving a *definition* of a unique valuation $\llbracket \cdot \rrbracket_I$ corresponding to interpretation I , but rather a set of *constraints* on legitimate valuation functions. In fact, we wish to keep these constraints quite weak, not insisting, for instance, that $\Phi \vee \Psi$ will be true if one of the disjuncts is true. This is in contrast with the “standard” definition of disjunction in situation semantics, which does make that stipulation (Barwise 1989; Fenstad 1987). So we are not insisting, for example, that a situation in which Mary sleeps must be one in which Mary sleeps or John sleeps. John may not be involved in the situation even to this extent. As a result our sentence intensions can be quite “fine-grained” – sufficiently so to avoid the usual difficulties with propositional attitudes such as belief in possible worlds semantics. Belief contexts apparently allow *some* substitutions (such as $\Phi \vee \Psi$ for $\Psi \vee \Phi$, and Φ for $\neg\neg\Phi$) but not ones that “introduce new material.”

As a final preliminary, we need to introduce the notion of the *parameters* of a formula, for use in our DRT-like treatment of indefinites (in generic conditionals and elsewhere). Intuitively, these are just the top-level existentially quantified variables of the formula (not embedded by quantifiers other than \exists or *The*). In essence we will treat \exists as a kind of “weak,” or “preemptable,” quantifier. It has its usual quantificational force only if the variable it quantifies is *a priori* valueless; otherwise, it is “ignored” (though it still serves to mark its variable *as* preemptable). The definite quantifier *The* will be treated similarly as preemptable. The *parameters* of formula Φ , written with an underscore function as $\underline{\Phi}$, are a set of variables defined as follows (where Q is a non-preemptable quantifier such as \forall , *Most*, ...):

$$(\exists \alpha: \Phi \Psi) = \{\alpha\} \cup \underline{\Phi} \cup \underline{\Psi}; \quad (\text{The } \alpha: \Phi \Psi) = \underline{\Phi} \cup \underline{\Psi};$$

$$[\underline{\Phi} \wedge \underline{\Psi}] = [\underline{\Phi} \vee \underline{\Psi}] = [\underline{\Phi} \text{ because } \underline{\Psi}] = \underline{\Phi} \cup \underline{\Psi};$$

$$\neg \underline{\Phi} = [\underline{\Phi} \rightarrow \underline{\Psi}] = [\underline{\Phi} \rightarrow_p \underline{\Psi}] = (Q\alpha: \Phi \Psi) = \emptyset;$$

$$\Box \underline{\Phi} = (\text{certainly } \underline{\Phi}) = \dots = \underline{\Phi};$$

$$[\underline{\Phi} * \eta] = [\underline{\Phi} ** \eta] = [\underline{\Phi} @ \eta] = \emptyset$$

If $d \in \mathcal{D}$, I is an interpretation of the atomic symbols of the logic, and α is a variable, $I(\alpha:d)$ denotes the interpretation identical with I except that it interprets α

as d (regardless of whether or not α already had a value under I). Also, if \underline{d} is a tuple of n elements of \mathcal{D} , and $\underline{\Phi}$ consists of n variables, then $I(\underline{\Phi};\underline{d})$ denotes the interpretation obtained from I by setting the denotations of those variables in $\underline{\Phi}$ which have no prior values to the corresponding individuals in \underline{d} (e.g., make the assignments in lexicographic order of the variables). In other words, if some of the variables in $\underline{\Phi}$ have no values under I and some do, then $I(\underline{\Phi};\underline{d})$ changes only the interpretations of the variables without prior values to the corresponding elements of \underline{d} , leaving denotations of variables with prior values unchanged.

The conditions satisfied by an interpretation I and any of its extensions to a valuation $\llbracket \cdot \rrbracket_I$ are as follows. Wherever $\llbracket \cdot \rrbracket$ occurs unsubscripted, it is an abbreviation for $\llbracket \cdot \rrbracket_I$. As before, we use $2 = \{0, 1\}$ as truth values and write $A \rightarrow B$ for the set of partial functions from A to B . $A^n \rightarrow B$ will abbreviate $(A \rightarrow (A \rightarrow \dots (A \rightarrow B) \dots))$.

1. <i>Type of atom α</i>	<i>$I(\alpha)$ is an element of:</i>
Individual constant:	\mathcal{D}
Individual variable:	\mathcal{D} , or is undefined ¹⁰
n -place predicate constant:	$\mathcal{D}^n \rightarrow (\mathcal{S} \rightarrow 2)$
Predicate modifier:	$(\mathcal{D} \rightarrow (\mathcal{S} \rightarrow 2)) \rightarrow (\mathcal{D} \rightarrow (\mathcal{S} \rightarrow 2))$
Sentence modifier:	$(\mathcal{S} \rightarrow 2) \rightarrow (\mathcal{S} \rightarrow 2)$
Predicate nominalization operator:	$(\mathcal{D} \rightarrow (\mathcal{S} \rightarrow 2)) \rightarrow \mathcal{D}$
Sentence nominalization operator:	$(\mathcal{S} \rightarrow 2) \rightarrow \mathcal{D}$

At this point it would be appropriate to provide characterizations of the specific nominalization operator *That* (denoting \wedge) and its inverse *True* (denoting \vee), as well as of the VP nominalization operator K_a and its inverse *Has-prop*. Since we take *That* to apply to *atemporal* sentences only (after deindexicalization), the propositions formed by \wedge presumably form a boolean algebra under analogs of the truth-functional connectives. The deindexicalization of infinitives (and gerunds) similarly involves introduction of explicit episodic terms, leading to a boolean algebra of properties. (These are properties of subject-episode *pairs*, not of subjects alone.) However, we do not yet understand the structure of propositions and properties well enough to elaborate these algebras and their relation to predicate and sentence intensions.

2. (*Upward persistence of predicates.*) If π is an n -place predicate constant, and s, s' are situations such that $s \preceq s'$, then $I(\pi)(d_1) \dots (d_n)(s) = I(\pi)(d_1) \dots (d_n)(s')$ for all $d_1, \dots, d_n \in \mathcal{D}$ such that the LHS is defined.
3. If α is an atomic expression, then $\llbracket \alpha \rrbracket = I(\alpha)$.
4. If π, α are expressions such that $\llbracket \pi \rrbracket \in A \rightarrow B$ and $\llbracket \alpha \rrbracket \in A$ for some sets A, B , then $\llbracket (\pi \alpha) \rrbracket = \llbracket \pi \rrbracket \llbracket \alpha \rrbracket$ (i.e., $\llbracket \pi \rrbracket (\llbracket \alpha \rrbracket)$). (More precisely, we ought to type functions

¹⁰Note that if a variable has a value, that value "preempts" \exists - and *The*-quantifiers for that variable.

syntactically, in parallel with the set-theoretic types in 1, and express the present rule in terms of the syntactic types of π and α , instead of their (in principle ambiguous) set memberships. But the intention should be clear.)

5. (*Upward persistence of formulas.*) If Φ is a formula, and s, s' are situations such that $s \preceq s'$, then $\llbracket \Phi \rrbracket^s = \llbracket \Phi \rrbracket^{s'}$ if the LHS is defined.
6. (*Persistence of atemporal formulas.*) If Φ is an atemporal formula, and s, s' are situations such that $s \sqsubseteq s'$, then $\llbracket \Phi \rrbracket^s = \llbracket \Phi \rrbracket^{s'}$ if the LHS is defined.

Atemporal formulas are those in which temporal predicates (such as *walk*, *girl*, *kiss*) occur only within the scope of $*$, $**$, or $@$. (Atemporal predicates such as $=$, *cause-of*, and *before* can occur anywhere in an atemporal formula.)

In the remaining clauses, Φ and Ψ are formulas, α is a variable, η is a term, and i is a time $\in \mathcal{I}$. Also, it will be convenient in the semantics of \exists and *The* to let $\mathcal{D}_{I(\alpha)}$ denote \mathcal{D} if $I(\alpha)$ is undefined, and singleton set $\{I(\alpha)\}$ otherwise. Further, we will use "restriction domains" $\mathcal{D}_{\alpha, \Phi, i} = \{d \mid \text{for some time } j \sqsubseteq i, \llbracket \Phi \rrbracket_{I(\alpha; d)}^j = 1\}$ in the semantics of quantifiers (clauses 10-13.) Finally, we say that a set of times \mathcal{J} *spans* a time i if $\min\{\text{begin}(j) \mid j \in \mathcal{J}\} = \text{begin}(i)$ and $\max\{\text{end}(j) \mid j \in \mathcal{J}\} = \text{end}(i)$. (For infinite \mathcal{J} , we use *glb* and *lub* instead of *min* and *max*.)

7. $\llbracket \neg \Phi \rrbracket^i = 1$ iff for all times $j \sqsubseteq i$, $\llbracket \Phi \rrbracket^j = 0$;
 $= 0$ iff for some time $j \sqsubseteq i$, $\llbracket \Phi \rrbracket^j = 1$.¹¹

E.g., $\neg[\text{John resign}]$ correctly describes a time interval just in case $[\text{John resign}]$ is false for each subinterval. This is not "outward persistent" for temporal Φ .

8. (a) *Symmetrical conjunction* (using $\underline{\sigma}$ for the parameter set $\Phi \wedge \Psi$):

$$\begin{aligned} \llbracket \Phi \wedge \Psi \rrbracket^i &= 1 \text{ iff for some } \underline{d} \in \mathcal{D}^{\underline{\sigma}} \text{ and some times } j, k \sqsubseteq i, \\ &\quad \llbracket \Phi \rrbracket_{I(\underline{\sigma}; \underline{d})}^j = \llbracket \Psi \rrbracket_{I(\underline{\sigma}; \underline{d})}^k = 1 \text{ and } \{j, k\} \text{ spans } i; \\ &= 0 \text{ iff for all } \underline{d} \in \mathcal{D}^{\underline{\sigma}} \text{ and all times } j, k \sqsubseteq i, \\ &\quad \text{either one of } \llbracket \Phi \rrbracket_{I(\underline{\sigma}; \underline{d})}^j, \llbracket \Psi \rrbracket_{I(\underline{\sigma}; \underline{d})}^k \text{ is } 0, \\ &\quad \text{or both are } 1 \text{ but } \{j, k\} \text{ does not span } i. \end{aligned}$$

- (b) *Asymmetrical conjunction* (using $\underline{\sigma}$ for the parameter set $\Phi \& \Psi$, i.e., $\Phi \cup \Psi$):

$$\begin{aligned} \llbracket \Phi \& \Psi \rrbracket^i &= 1 \text{ iff for some } \underline{d} \in \mathcal{D}^{\underline{\sigma}}, \llbracket \Phi \rrbracket_{I(\underline{\sigma}; \underline{d})}^{\text{init}(i)} = \llbracket \Psi \rrbracket_{I(\underline{\sigma}; \underline{d})}^i = 1; \\ &= 0 \text{ iff for all } \underline{d} \in \mathcal{D}^{\underline{\sigma}}, \llbracket \Phi \rrbracket_{I(\underline{\sigma}; \underline{d})}^{\text{init}(i)} = 0, \text{ or } \llbracket \Psi \rrbracket_{I(\underline{\sigma}; \underline{d})}^i = 0. \end{aligned}$$

¹¹In general, we will not be able to say, even for times i , that $\llbracket \Phi \rrbracket^i = 0$ wherever $\llbracket \Phi \rrbracket^i \neq 1$, because the expressiveness of the logic inevitably leads to truth value gaps (or inconsistency). For instance, we can express the Russell sentence $[(K_\lambda \lambda p \neg[p \text{ Has-prop } p]) \text{ Has-prop } (K_\lambda \lambda p [p \text{ Has-prop } p])]$ ("The property of not being self-applicable has the property of being self-applicable").

Remarks: Asymmetrical conjunction is useful for abbreviating the semantics of quantification. Both (a) and (b) allow for forward and backward anaphora, through the "parameter" mechanism. In effect, they are evaluated as if existential quantifiers at the highest level in Φ and Ψ had wide scope over the entire conjunction. (So, "Some man x is ill, and x is coughing" is evaluated as "For some man x , x is ill and x is coughing.")¹²

9. $\llbracket \Phi \vee \Psi \rrbracket^i = 1$ iff $\llbracket \Phi \rrbracket^i = 1$ or $\llbracket \Psi \rrbracket^i = 1$;
 $= 0$ iff $\llbracket \Phi \rrbracket^i = 0$ and $\llbracket \Psi \rrbracket^i = 0$.
10. $\llbracket (\exists \alpha: \Phi \Psi) \rrbracket^i = 1$ iff for some $d \in \mathcal{D}_{I(\alpha)}$, $\llbracket \Phi \& \Psi \rrbracket_{I(\alpha:d)}^i = 1$;¹³
 $= 0$ iff for all $d \in \mathcal{D}_{I(\alpha)}$, $\llbracket \Phi \& \Psi \rrbracket_{I(\alpha:d)}^i = 0$.

Remarks: Our use of asymmetrical conjunction here reflects our observation that indefinites seem to require truth of the restriction predicate at least at the *beginning* of the episode described, but not necessarily at the end. Examples are "A child grew up," "A bubble burst," "An ice cube melted," etc. Concerning examples like "A planet formed," "He became an adult," or "Mary baked a cake," we feel that these require an intensional account, i.e., the verbs are essentially predicate operators. Note the contrasting valid and invalid inferences from the progressive forms: A child was growing up, therefore there was a child; A planet was forming, *therefore there was a planet; etc.

For sentences like "A little boy (who used to live here) is now a man," and "A certain violinist was once a child prodigy" (cf. Enç 1981, Hinrichs 1988), we conjecture that there are two additional phenomena involved: (i) A nominal (predicate) operator *sometime*, where $[\alpha \text{ (sometime } \pi)]$ means that " α is at some time π ," i.e., it will be true at *all* times, if it is ever true (so the result is *atemporal*). This appears to be involved in the first sentence, and in most sentences with *event*-nominals, such as "A supernova was observed at Palomar." (ii) A scope phenomenon, in which a quantifier escapes from the scope of the tense operator, thereby getting its temporal reference from the utterance event; this seems to be involved in one reading of *a certain violinist* in the second sentence, viz., the reading where this has *present* reference. (In another reading, (*sometime violinist*) is involved.)

11. $\llbracket (\text{The } \alpha: \Phi \Psi) \rrbracket^i = 1$ iff for the unique $d \in \mathcal{D}_{I(\alpha)}$ satisfying $\llbracket \Phi \rrbracket_{I(\alpha:d)}^{init(i)} = 1$,
 $\llbracket \Phi \& \Psi \rrbracket_{I(\alpha:d)}^i = 1$;

¹² Note that according to these truth conditions, names of existentially quantified variables matter. For example, $\llbracket (\exists x: \Phi \Psi) \wedge (\exists x: \Phi \neg \Psi) \rrbracket$ will be logically false, since there is only one parameter, x , which is varied "simultaneously" in both conjuncts. On the other hand, $\llbracket (\exists x: \Phi \neg \Psi) \wedge (\exists y: \Phi \neg \Psi) \rrbracket$ has the usual truth conditions (unless x, y have prior values).

¹³ Note that by the definition of $\mathcal{D}_{I(\alpha)}$ this is equivalent to: iff $I(\alpha)$ is defined, and $\llbracket \Phi \wedge \Psi \rrbracket^i = 1$, or $I(\alpha)$ is undefined, and for some $d \in \mathcal{D}$, $\llbracket \Phi \rrbracket_{I(\alpha:d)}^i = 1$.

$$= 0 \text{ iff for the unique } d \in \mathcal{D}_{I(\alpha)} \text{ satisfying } \llbracket \Phi \rrbracket_{I(\alpha:d)}^{init(i)} = 1, \\ \llbracket \Phi \& \Psi \rrbracket_{I(\alpha:d)}^i = 0.$$

Remarks: This leads to a “referential” reading if α has a prior value, and a situation-dependent “Russellian” one otherwise, except that we have deliberately left open the question of whether non-uniqueness should lead to falsity or undefinedness.

12. $\llbracket (\forall \alpha: \Phi \Psi) \rrbracket^i = 1$ iff there is a total function $f: \mathcal{D}_{\alpha, \Phi, i} \rightarrow \mathcal{I}$ such that
 for all $d \in \mathcal{D}_{\alpha, \Phi, i}$: $f(d) \subseteq i$, $\llbracket \Phi \& \Psi \rrbracket_{I(\alpha:d)}^{f(d)} = 1$, and
 $range(f)$ spans i ;
 = 0 iff for all total functions $f: \mathcal{D}_{\alpha, \Phi, i} \rightarrow \mathcal{I}$ such that
 for all $d \in \mathcal{D}_{\alpha, \Phi, i}$, $f(d) \subseteq i$:
 either for some $d \in \mathcal{D}_{\alpha, \Phi, i}$, $\llbracket \Phi \& \Psi \rrbracket_{I(\alpha:d)}^{f(d)} = 0$,
 or $range(f)$ does not span i .

Remarks: Note that the episodes quantified over may be arbitrarily dispersed in time, as in “Every premature baby was put in an incubator.” However, the restriction predicate is assumed to apply at the beginning of each component event, much as in 11. (Again, *sometime* and scoping phenomena can lead to apparent violations of this constraint.) The use of a function in the definition is aimed at the following kind of problem. Suppose that there are three relevant men, Tom, Dick, and Harry; Tom gives Mary twelve roses, one at a time; then Dick and Harry each give her a rose; then Tom gives Mary a dozen more, again, one at a time. Can this entire event be characterized as *each man giving Mary a rose*? We think not, though it *contains* two such events (one with Tom first, the other with Tom last).

Also, note that as it stands the condition for truth value 1 requires a nonempty domain of restricted quantification, $\mathcal{D}_{\alpha, \Phi, i}$, since the empty set cannot span any time interval. This could easily be adjusted to give truth value 1 for empty $\mathcal{D}_{\alpha, \Phi, i}$.

13. $\llbracket (\text{Most } \alpha: \Phi \Psi) \rrbracket^i = 1$ iff there is a partial function $f: \mathcal{D}_{\alpha, \Phi, i} \rightarrow \mathcal{I}$ such that
 $dom(f)$ includes more than half¹⁴ of $\mathcal{D}_{\alpha, \Phi, i}$, and
 for all $d \in dom(f)$, $f(d) \subseteq i$ and $\llbracket \Phi \& \Psi \rrbracket_{I(\alpha:d)}^{f(d)} = 1$, and
 $range(f)$ spans i ;
 = 0 iff for all partial functions $f: \mathcal{D}_{\alpha, \Phi, i} \rightarrow \mathcal{I}$ such that
 $dom(f)$ includes at most half of $\mathcal{D}_{\alpha, \Phi, i}$, and
 for all $d \in \mathcal{D}_{\alpha, \Phi, i}$, $f(d) \subseteq i$:
 either for all $d \in dom(f)$, $\llbracket \Phi \& \Psi \rrbracket_{I(\alpha:d)}^{f(d)} = 0$,
 or $range(f)$ does not span i .

¹⁴ For infinite domains, we hope eventually to use a distributional approach, similar to that of Bacchus (1988a,b).

Remarks: Conditions similar to 12 and 13 can be formulated for other *monotone increasing* quantifiers, such as *many*, *at least one hundred*, etc. These are quantifiers Q such that for atemporal Φ , Ψ and Υ , $(Q\alpha:\Phi\Psi)$ and $(\forall\alpha:\Psi\Upsilon)$ together entail $(Q\alpha:\Phi\Upsilon)$ (see Barwise and Cooper 1981). We assume that non-monotone increasing quantifiers can be recast in terms of monotone increasing ones and negation. Thus, $(\text{Few}\alpha:\Phi\Psi)$ becomes $\neg(\text{Many}\alpha:\Phi\Psi)$, and $(\text{Exactly-}n\alpha:\Phi\Psi)$ becomes $(\text{At-least-}n\alpha:\Phi\Psi) \wedge \neg(\text{More-than-}n\alpha:\Phi\Psi)$, etc.

14. $\llbracket \Phi \rightarrow \Psi \rrbracket^i = 1$ iff $\llbracket \Phi \rrbracket^j = 0$ for all times $j \sqsubseteq i$, or $\llbracket \Phi \wedge \Psi \rrbracket^i = 1$;
 $= 0$ iff $\llbracket \Phi \rrbracket^j = 1$ for some time $j \sqsubseteq i$, and $\llbracket \Phi \wedge \Psi \rrbracket^i = 0$.
 For s any situation, $\llbracket \Phi \rightarrow \Psi \rrbracket^s = 0$ if $\llbracket \Phi \rrbracket^s = 1$ and $\llbracket \Phi \wedge \Psi \rrbracket^s$ is undefined.

Remarks: Through 8(a), this plausibly handles many nongeneric “donkey sentences,” such as “If Pedro owns a donkey, he will ride it to town tomorrow” (*cf.*, Schubert and Pelletier 1989). This will *not* require Pedro to ride *all* his donkeys to town.

The third condition ensures that $\Phi \rightarrow \Psi$ always has a truth value if its antecedent is true. This strengthens the conditional sufficiently so that it can be used together with necessity to express meaning postulates (analogous to *constraints* in situation semantics), guaranteeing that if a situation contains the “antecedent information,” it also contains the “consequent information.”

15. $\llbracket \Box \Phi \rrbracket^i = 1$ iff for all times $j \in \mathcal{I}$, $\llbracket \Phi \rrbracket^j = 1$;
 $= 0$ iff for some time $j \in \mathcal{I}$, $\llbracket \Phi \rrbracket^j = 0$.
16. $\llbracket \Phi \rightarrow_p \Psi \rrbracket^i = 1$ iff for “at least a proportion p ” of elements
 $\underline{d} \in \mathcal{D}[\Phi]$ such that $\llbracket \Phi \rrbracket_{I(\underline{\Phi}:\underline{d})}^i = 1$,
 $\llbracket \Phi \wedge \Psi \rrbracket_{I(\underline{\Phi}:\underline{d})}^i = 1$;
 $= 0$ iff for “more than a proportion $(1-p)$ ” of elements
 $\underline{d} \in \mathcal{D}[\Phi]$ such that $\llbracket \Phi \rrbracket_{I(\underline{\Phi}:\underline{d})}^i = 1$,
 $\llbracket \Phi \wedge \Psi \rrbracket_{I(\underline{\Phi}:\underline{d})}^i = 0$.

Remarks: This is only a rough approximation to what is required. First, the talk of “proportions” needs to be replaced by a notion of measure, based on a distribution over individuals. Second, instead of using purely “extensional statistics” (proportions of Ψ -instances relative to Φ -instances at i), the truth conditions should in general be modalized to reflect the *normic* character of many generic sentences; this could be done by evaluating the antecedent and consequent not only at i , but also at “nearby” *Nonactual* times extending over the same clock-time range as i . Third, the definition as it stands does not properly accommodate backward anaphora, as in “If *he* owns a donkey, a *man* often beats it.” There are also some subtle problems that cannot be discussed here concerning which parameters (indefinites) should be “iterated” over, and which ones should not.

Nevertheless, the rule is sufficiently accurate to be a guide to action. Given that Φ holds at a time for certain values of its parameters, it is reasonable on the basis of much-discussed inductive principles (e.g., Kyburg 1983, Bacchus 1988a,b) to assign conclusion Ψ , for those values, degree of certainty $\geq p$ (at least in the absence of other information).

17. $\llbracket \Phi * \eta \rrbracket^i = 1$ iff $\llbracket \eta \rrbracket \subseteq \text{world}(i)$ and $\llbracket \Phi \rrbracket^{\llbracket \eta \rrbracket} = 1$;
 $= 0$ iff $\llbracket \eta \rrbracket \not\subseteq \text{world}(i)$ or $\llbracket \Phi \rrbracket^{\llbracket \eta \rrbracket} \neq 1$.

Remark: Thus $*$ denotes truth over an (actual) episode.

18. $\llbracket \Phi @ \eta \rrbracket^i = 1$ iff Φ is stative and $\llbracket \Phi * \eta \rrbracket^i = 1$, or
 Φ is nonstative and there is an $s \subseteq \llbracket \eta \rrbracket$ such that
 $\llbracket \Phi * e \rrbracket_{(e:s)}^i = 1$, where e is a variable not occurring in Φ ;
 $= 0$ iff Φ is stative and $\llbracket \Phi * \eta \rrbracket^i = 0$, or
 Φ is nonstative and there is no $s \subseteq \llbracket \eta \rrbracket$ such that
 $\llbracket \Phi * e \rrbracket_{(e:s)}^i = 1$, where e is a variable not occurring in Φ .

One can also define the "@" connective syntactically as follows:

$$\llbracket \Phi @ \eta \rrbracket =_{df} \begin{cases} \llbracket \Phi * \eta \rrbracket & \text{if } \Phi \text{ is stative;} \\ (\exists e: [e \text{ subep } \eta] \llbracket \Phi * e \rrbracket) & \text{if } \Phi \text{ is nonstative,} \end{cases}$$

where *subep* corresponds to the \subseteq relation, i.e., $\llbracket \text{subep} \rrbracket \in \mathcal{D}^2 \rightarrow \mathcal{S} \rightarrow 2$ such that for all $s_1, s_2 \in \mathcal{S}$ and $w \in \mathcal{W}$, $\llbracket \text{subep} \rrbracket^{s_1 s_2 w} = 1$ iff $s_2 \subseteq s_1$.

Remark: The stative/nonstative distinction plays an important role in most theories of aspect and aspectual class (e.g., Dowty 1979,86, ter Meulen 88). We make use of the distinction in the interpretation of perfective aspect (see next section) and in certain meaning postulates (see footnote 29) and plausible inference rules (similar to "causal connection" in section 5). Our aspectual classes are quite rudimentary compared, for instance, to Moens and Steedman's (1988) who have a very detailed (but informal) analysis of the interaction between temporal reference and aspectual classes. One test for stativity is compatibility with durative adverbials (*was ill for two days, waited for an hour* versus **resigned for two minutes, *planted a tree for an hour*). Note that on this test negative VPs (such as *has not planted a tree (for an hour)*) are stative.

19. $\llbracket \Phi ** \eta \rrbracket^i = 1$ iff $\llbracket \Phi * \eta \rrbracket^i = 1$, and
for all propositions $p \in \mathcal{P}$ such that $\vee p(\llbracket \eta \rrbracket) = 1$, and
for all $s \in \mathcal{S}$, if $\llbracket \Phi \rrbracket^s = 1$ then $\vee p(s) = 1$;
 $= 0$ iff $\llbracket \Phi * \eta \rrbracket^i = 0$, or
for some proposition $p \in \mathcal{P}$ such that $\vee p(\llbracket \eta \rrbracket) = 1$, and
for some $s \in \mathcal{S}$, $\llbracket \Phi \rrbracket^s = 1$ and $\vee p(s) \neq 1$.

Remark: This expresses that if η is completely described by Φ , then the only additional properties it has are entailments of Φ .

20. For π a formula or n -place predicative expression ($n \geq 1$),
 $\llbracket \lambda \alpha \pi \rrbracket = \{ \langle d, \llbracket \pi \rrbracket_{\alpha:d} \rangle \mid d \in \mathcal{D}, \llbracket \pi \rrbracket_{\alpha:d} \text{ defined} \}.$
21. For τ_1, τ_2 terms,
 $\llbracket \tau_1 = \tau_2 \rrbracket^i = 1$ iff $\llbracket \tau_1 \rrbracket = \llbracket \tau_2 \rrbracket$, with both defined;
 $= 0$ iff $\llbracket \tau_1 \rrbracket \neq \llbracket \tau_2 \rrbracket$, with both defined.

This completes our enumeration of truth conditions. Primary omissions are a detailed algebraic specification of the various nominalization operators and their inverses, and the semantics of all but a few quantifiers (which can be “filled in” more or less analogously). Also, we have not specified a semantics for the important sentential connective *because*, about which we know little more than we said in section 2. (There are as well additional operators in our logic relating to questions and *wh*-nominals which we have not mentioned, for lack of semantic details). Finally, we have not given any semantics of functions over individuals though clauses 1-4 might easily have been formulated to cover them. The reason is that we are not convinced that any words or phrases of English express functions from individuals to individuals. One commonly sees such relations as *father of* treated functionally in NL systems, but it seems to us that it is a matter of contingent world knowledge, not a matter of logic, that fathers are unique (and genetic advances might falsify this). Even such relations as *weight of* seem to us only contingently functional, and more importantly, language appears not to distinguish apparent functions like *weight of* from evident relations like *neighbor of* or *advantage of*. One function we do use is the pairing function, whose semantics is simply $\llbracket [\alpha \mid \beta] \rrbracket = \langle \llbracket \alpha \rrbracket, \llbracket \beta \rrbracket \rangle$ (an ordered pair which is defined whenever both elements are defined). Also, we did mention arithmetic functions at one point, which we find convenient, but they have little bearing on our concerns in this paper.

Another postscript to our truth conditions concerns the fact that our semantics of constants in clause 1 makes them rigid designators, denoting specific individuals independently of situations. This in turn leads to rigid identity conditions, as in the final clause 21. The rigidity assumption keeps truth conditions for predication and quantification relatively simple, and straightforwardly accounts for *de re* readings in modal contexts. On the other hand, it invites trouble with *de dicto* readings such as are needed to resolve the Morning Star - Evening Star puzzle.

A nice example from our target story, *Little Red Riding Hood*, is the circumstance that the heroine (let us call her LRRH) thought the wolf, in Grandmother's clothes and lying in her bed, was Grandmother. This is not representable as

$$\llbracket \text{LRRH believe (That } [\llbracket W = G \rrbracket * E1]) \rrbracket ** E2 \rrbracket,$$

for some episodes $E1, E2$. LRRH certainly did not think W (the wolf) and G (Grandmother) were one and the same individual. Rather, we need a *de re* reading for W and a *de dicto* reading for G , i.e., LRRH believed *of* W that he was G .

One possibility would be to associate an individual concept $c_\alpha \in \mathcal{S} \rightarrow \mathcal{D}$ with each constant α , and in effect to “scope” constants using a quantifier-like (*Load* $\alpha \Phi$) con-

struct to provide a local value for α via c_α . For instance, instead of the above we would have

$$[[\text{LRRH believe (That (Load G [[W = G] * E1])}] ** E2],$$

where for a top-level formula, all constants are assumed to be initially *Load*-ed with $c_\alpha(i)$ for the time of evaluation i . More deeply embedded "*Load* α "s will override those with wider scope, for the same α . The semantics of *Load*-ing would be

$$[[\text{Load } \alpha \Phi]]_I^s = [[\Phi]]_{I(\alpha: c_\alpha(s))}^s.$$

Note that apart from changes in clause 1, other truth conditions would not be affected since constants in such a scheme are still interpreted as individuals, though now they may be different individuals in different situations.¹⁵

Entailment in our logic is defined for a set of premises Σ and a conclusion Ψ as follows. We use the notion of an *MP-constrained* valuation function $[[\cdot]]$, i.e., one which conforms with the constraints we have stated and also satisfies $[[\Phi]]^w = 1$ for all worlds $w \in \mathcal{W}$ and all formulas Φ designated as meaning postulates. Then

$$\begin{aligned} \Sigma \models \Psi \text{ iff for all MP-constrained valuation functions } [[\cdot]] \\ \text{and all worlds } w \in \mathcal{W} \text{ such that } [[\Phi]]^w = 1 \text{ for all } \Phi \in \Sigma, \\ [[\Psi]]^w = 1. \end{aligned}$$

Note that by defining entailment for worlds, it automatically applies only to atemporal formulas (since only atemporal ones can be true in a world, i.e., for all time).

We will prove the following entailment, as an illustration of how the truth conditions for times interact with the persistence conditions (5) and (6):

$$\begin{aligned} \text{For } \Phi \text{ a formula and } \eta \text{ a term,} \\ [(\neg\Phi) * \eta] \models (\forall t: [[t \text{ time}] \wedge [t \text{ during } \eta]] (\neg[\Phi * t])), \end{aligned}$$

with $I(\text{time}) \in \mathcal{D} \rightarrow \mathcal{S} \rightarrow 2$ such that for all $s \in \mathcal{S}$ and $i \in \mathcal{I}$, $I(\text{time})(s)(i) = 1$ if $s \in \mathcal{I}$ and $s \sqsubseteq i$, and $= 0$ otherwise (so that only *Actual* times satisfy the predicate); and $I(\text{during}) \in \mathcal{D}^2 \rightarrow \mathcal{S} \rightarrow 2$ such that for all $s_1, s_2 \in \mathcal{S}$ and $i \in \mathcal{I}$, $I(\text{during})(s_1)(s_2)(i) = 1$ if $\text{begin}(s_1) \leq \text{begin}(s_2)$ and $\text{end}(s_2) \leq \text{end}(s_1)$, and $= 0$ otherwise.

Proof. For any world $w \in \mathcal{W}$, $[[(\neg\Phi) * \eta]]^w = 1$ implies $[[\eta]] \sqsubseteq w$ and $[[\neg\Phi]]^w = 1$ (by 17). Hence if i is the time concurrent with $[[\eta]]$ in w , $[[\neg\Phi]]^i = 1$ (by upward persistence, 5). Hence, for all times $j \sqsubseteq i$, $[[\Phi]]^j = 0$ (by 7). Hence for all times $j \sqsubseteq i$, $[[\Phi * t]]_{I(t:j)}^j = 0$ (by 17). Hence for all times $j \sqsubseteq i$, $[[\neg[\Phi * t]]]_{I(t:j)}^i = 1$ (by

¹⁵ A plausible way of introducing *Load* in the representation of a sentence like the above may be through a semantic rule that interprets a name N serving as complement of a copular verb as $\lambda x(\text{Load } N'[x = N'])$, where N' is the logical translation of N . This would explain why "LRRH believed that the wolf was Grandmother" is not synonymous with "LRRH believed that Grandmother was the wolf." Note that *Load* must be assumed to block substitution for free variables in its scope, if the substituted values involve the "*Loaded*" constant (at least in intensional contexts).

7). Hence for all times $j \sqsubseteq i$, $\llbracket \neg[\Phi * t] \rrbracket_{I(t:j)}^w = 1$ (by persistence, 6). Then by the constraints placed above on $I(\text{time})$ and $I(\text{during})$, and the truth conditions for “ \forall ” and “ \wedge ,” $\llbracket (\forall t: \llbracket [t \text{ time}] \wedge [t \text{ during } \eta] \rrbracket (\neg[\Phi * t])) \rrbracket^w = 1$. \square

4 From English to Episodic Logic

It is one thing to posit logical forms for particular sentences, but quite another to generate them systematically through a grammatical/semantic formalism. An important advantage of our representation is that it can be directly and uniformly computed from syntactic analyses of input sentences.

We conceive of this computation as involving three processing phases (to be interleaved eventually). First, we use a GPSG-style grammar to compute indexical (context-dependent) translations with ambiguously scoped quantifiers, connectives, and tense operators, such as

$[\langle \exists \text{ person} \rangle \langle \text{past kiss} \rangle \text{ Mary}]$

(“Someone kissed Mary”). Second, we scope the unscoped operators, obtaining formulas such as

$(\text{past } (\exists x: [x \text{ person}] [x \text{ kiss Mary}]))$.

This is still indexical in that it is past *relative* to some implicit utterance time, and perhaps also involves implicit relations between the described event and previously described events. Third, we combine the indexical translation with a *context structure* (whose most important component, for present purposes, will be a stack of temporal structures called *tense trees*), and then apply equivalence transformations to the combination, which recursively eliminate the dependence on context, ultimately giving the desired nonindexical (context-independent) translation.¹⁶

We will illustrate the derivation of a logical form for the sentence

(14) John inferred that Mary was tired.

Here is a GPSG fragment adequate for (14), where each lexical or phrase structure rule is paired with a corresponding semantic rule:

- A. $\text{NP} \rightarrow \text{Mary}; \text{Mary}$
- B. $\text{NP} \rightarrow \text{John}; \text{John}$
- C. $\text{A}[\text{pred}] \rightarrow \text{tired}; \text{tired}$
- D. $\text{AP} \rightarrow \text{A}; \text{A}'$

¹⁶In the *KR'89* paper, we tried to formulate the semantic rules of the GPSG grammar so that they immediately introduce episodic variables and “*” and “**” operators, and express tense, aspect and adverbials in terms of relations between episodes. However, this ran into difficulties with the interaction between tense, perfective aspect, negation and time adverbials. The unnatural *t*, *e*, and *h* features in our previous fragment are symptoms of some of these difficulties.

- E. $V[\text{be, past, 3per, sing, sub}] \rightarrow \text{was}; \lambda P <\text{subpast } P>$ ¹⁷
- F. $V[\text{be, past, 3per, sing, sub}] \rightarrow \text{was}; \lambda P <\text{past } P>$
- G. $VP \rightarrow V[\text{be}] AP[\text{pred}]; (V' AP')$
- H. $S \rightarrow NP VP; [NP' VP']$
- I. $S[\text{that, tense}^{18}] \rightarrow \text{COMPL}[\text{that}] S[\text{sub, tense}]; (\text{That } S')$
- J. $V[_S[\text{that}], \text{past}] \rightarrow \text{inferred}; <\text{past infer}>$
- K. $VP \rightarrow V[_S[\text{that}]] S[\text{that}]; (V' S')$

Certain feature principles are assumed here – namely, certain versions of the head feature principle, the control agreement principle, and the subcategorization principle (cf. Gazdar *et al.* (1985) and Pollard and Sag (1987)).¹⁹

The initial translations computed for (14) are then

- (14' a) [John $<\text{past infer}>$ (That [Mary $<\text{subpast tired}>$])], and
- (14' b) [John $<\text{past infer}>$ (That [Mary $<\text{past tired}>$])].

Note that we get two translations since rules E and F provide alternative semantic formulas for past-tensed *be* in embedded sentences. We will pursue only (14' a).

The second processing phase here consists of “raising” the occurrences of *past* and *subpast* to a permissible sentential level. (We follow, e.g., Richards and Heny (1982) in treating tense as sentence operators with wide scope over adverbials. This is in contrast with Hinrichs (1988) who gives tense scope only over the predicate that corresponds to the main verb.) This has the following unique result, since *That* acts as “scope trap”:

- (14'') (past [John infer (That (subpast [Mary tired]))]).

In the third phase, we need to combine this with a context structure for the utterance and “deindexicalize” it. This involves application of a fixed set of equivalences called *deindexicalization rules*, of which we will show only a few, as needed. A context structure may in general contain such objects as a *history list* (with recent referents), a *focus list* (with recently focused entities), and other contextual information (see Allen 1987). However, as mentioned, the only part of the context structure we will be concerned with here is the stack of *tense trees*, specifically, the *current* (most recent) tense tree. A node of a tense tree may have up to three branches, where leftward branches correspond to *past* tense (going to an earlier episode), straight downward branches to *perfective aspect* (going to an earlier episode or one just ending), and rightward branches to *future*


¹⁷ “sub” is a head feature that indicates that the constituent appears in a verb complement clause (e.g., the object of a propositional attitude); “subpast” means semantic subordinate past. Note that not every syntactically subordinate past is semantically subordinate past.

¹⁸ Rather than treating features as attribute-value functions, as in standard GPSG, we treat them as trees. In this instance, *tense* is the root of a feature tree with daughters *pres* and *past*. A tree node is considered compatible with any of its ancestors or descendants.

¹⁹ The subcategorization principle obviates the need for explicit rules like K, but we show the rule for greater clarity. The departures from the usual formulations are not important for present purposes.


modality (going to a later episode). Each node contains a list of (recently mentioned) episodes. A tense tree always has exactly one node in focus, marked as \odot ; and processing of a top-level sentence always begins with the root node in focus and a token for the utterance of the sentence as the rightmost (most recently added) element of its episode list.

Let us assume that (14) was preceded by another sentence such as "Mary looked pale." Then the initial tense tree for (14) might be



where $u0$ and $u1$ denote the utterances of the preceding sentence and sentence (14), respectively, and $e0$ denotes the episode reported in the preceding sentence (Mary's looking pale).

Our task is to deindexicalize (14") using a context structure containing this tree as its current tense tree. We normally write an indexical formula Φ in combination with a context structure C as Φ_C , but since we are neglecting all but the current tree T here, we shall instead write it as Φ_T . For (14") and the assumed T , this is



The first relevant deindexicalization rule is for deindexicalizing *past*:

$$\text{Past: } (\text{past } \Phi)_T \leftrightarrow (\exists e_T: [[e_T \text{ before Ref}_T] \wedge [e_T \text{ successor-ep Ref}_{\swarrow T}]] [\Phi_{O\swarrow T} ** e_T]).$$

Here Ref_T denotes the most recently stored episode at the focal node of T . e_T is the "next episode variable not yet used in T ." This is assumed to be uniquely defined for a given T ²⁰; for instance, it might be the letter e followed by the least numeric suffix i such that ei does not occur anywhere in T . " \swarrow " (called "shift left") shifts the focus to the left daughter, i.e., to the past, creating a new daughter if none yet exists. Note that the rule explicitly records the fact that e_T appeared in the linguistic context as successor episode to $\text{Ref}_{\swarrow T}$ (which denotes the most recently stored episode at the focal node of $\swarrow T$, i.e., in our example, $e0$).²¹ This *successor-ep* predication can later be used to make *probabilistic narrative inferences* about the temporal or causal relations between the two episodes (see section 5). Next, " \circ " (called "store") is a function that operates on a tense tree T , placing a new episodic variable e_T at the focal node; i.e., $\text{Ref}_{O\swarrow T} = e_T$. Thus, $O\swarrow T$, that is, $(O(\swarrow T))$, denotes the transformed tense tree obtained from T by first shifting the focus down to the left, and then storing the new episode e_T at the new

²⁰ Actually, for a given C , but we are neglecting the rest of C here.

²¹ If $\text{Ref}_{\swarrow T}$ is undefined, the *successor-ep* predication is omitted.

focal node.

Applying the *Past*-rule to (15), we obtain

$$(16) (\exists e1: [e1 \text{ before } u1] \wedge [e1 \text{ successor-ep } e0]) \\ [[\text{John infer (That (subpast [Mary tired]))}]] \begin{array}{c} \nearrow u0, u1 \\ \odot e0, e1 \end{array} **e1])$$

Note that both the translation and the focus of the tense tree now contain the token *e1* for John's "inferring" episode. Since the new tense tree is associated with the predication [John infer (That ...)], a predication deindexicalization rule (which supplies the tense tree to the last argument) is applied:

$$\text{Pred: For } \pi \text{ an atomic predicate and } \tau_1, \dots, \tau_n \text{ terms, where } \tau_i, 1 \leq i \leq n, \\ \text{is atomic, except possibly for } \tau_n:^{22} \\ [[\tau_1 \pi \tau_2 \dots \tau_n]_T ** \eta] \leftrightarrow [[\tau_1 \pi \tau_2 \dots \tau_{n-1} \tau_n] ** \eta]$$

The result of applying *Pred* to (16) is

$$(17) (\exists e1: [e1 \text{ before } u1] \wedge [e1 \text{ successor-ep } e0]) \\ [[\text{John infer (That (subpast [Mary tired]))}]] \begin{array}{c} \nearrow u0, u1 \\ \odot e0, e1 \end{array}] **e1)$$

Next, the deindexicalization of the term (*That ...*), with the same tense tree, generates a new single-node tree "embedded" by the given tree (note the horizontal "embedding link"):

$$(18) \dots (\text{That (subpast [Mary tired])} \begin{array}{c} \nearrow u0, u1 \\ \odot e1 \end{array}) \dots$$

This is done by a *That*-rule, whose details we omit. Note that the root of the new tree becomes the new focus, and that the previous *Ref_T*, i.e., *e1*, has been copied to that focus. Though we are now strictly dealing with *two* trees connected by an embedding link, we shall still refer to them jointly as "the" tense tree.

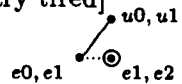
Next, one of the rules for *Subpast* is applied:

$$\text{Subpast(1): For } T \text{ focused at a past-dominated node, i.e., one with a leftward} \\ \text{branch in its ancestry (where ancestry also passes through embed-} \\ \text{ding links):} \\ (\text{subpast } \Phi)_T \leftrightarrow (\exists e_T: [e_T \text{ ends-during } \text{Ref}_T] [\Phi_{OT} ** e_T])$$

Note that this interprets the *past*-embedded episode as *overlapping* the embedding

²² A slightly more complex rule, allowing successive arguments to change the context, is needed for the general case of nonatomic arguments.

episode. (*Subpast*(2) is for T focused at a non-*past-dominated* node, which interprets the *past*-embedded episode as occurring *before* the embedding episode.) For the formula (*subpast* ...) within (18), the result is

$$(19) \dots (\exists e2:[e2 \text{ ends-during } e1] [[\text{Mary tired}]_{u0, u1} **e2]) \dots$$


Note that though *e2* (Mary's being tired) is predicated to end during *e1* (John's inferring it), it may well be an initial segment of a much longer episode of the same type; i.e., Mary may continue to be tired.²³ It remains to apply the *Pred*-rule once more, and to interpret *Mary* relative to the tense tree as just *Mary* (a nonindexical constant), with overall result:

$$(20) (\exists e1:[e1 \text{ before } u1] \wedge [e1 \text{ successor-ep } e0]) \\ [[\text{John infer (That } (\exists e2:[e2 \text{ ends-during } e1] [[\text{Mary tired}] **e2]) **e1)]$$

Thus we have a fully context-independent representation of (11), which can be used freely for inference (including narrative inferences based on the *successor-ep* relation).

There are some eight rules altogether for deindexicalizing tense, future modality, and perfective and progressive aspect, and additional rules to deal with other constructs of our logic. We must leave most details to another paper oriented toward these linguistic issues (Schubert and Hwang, to appear), but will briefly consider perfective aspect, conjunctions, quantification, and, at slightly greater length, adverbials.

Perfective

Syntactically, the perfective is handled straightforwardly through the auxiliary VP rule,

$VP \rightarrow V[\text{aux}] \text{ ADVL}[\text{pre-VP}]^- \text{ ADV}[\text{not}]^- \text{ XP}; (\text{ADVL}' \lambda x(\text{ADV}' [x (V' \text{ XP}')]))$,²⁴ and lexical rules such as

$$\begin{aligned} V[\text{aux}, \text{pres}, 3\text{per}, \text{sing}, \text{VP}[-\text{en}]] &\rightarrow \text{has}; \lambda P \lambda x < \text{pres} (\text{perf}_1 [x P]) >, \\ V[\text{aux}, \text{past}, \text{pers}, \text{numb}, \text{VP}[-\text{en}]] &\rightarrow \text{had}; \lambda P \lambda x < \text{past} (\text{perf}_1 [x P]) >, \\ V[\text{aux}, \text{past}, \text{pers}, \text{numb}, \text{VP}[-\text{en}]] &\rightarrow \text{had}; \lambda P \lambda x < \text{past} (\text{perf}_2 [x P]) >. \end{aligned}$$

Note that we distinguish two perfective operators in the semantics, *perf*₁ and *perf*₂. As will be seen, the interpretation of *perf*₁ is sensitive to the stative/nonstative distinction,

²³On the other hand, if the embedded episode were *non-stative*, as in "Mary noticed that John winked at her," it would indeed be prevented from extending beyond the embedding episode.

²⁴"minus"-superscripts indicate optional constituents; where such constituents are not present, the correct semantic rule is obtained by replacing their translations by the identity operator, λPP . The rule is intended to cover all auxiliaries preceding the main verb, as well as copular *be*, i.e., $V[\text{be}, \text{XP}[\text{pred}]]$ (so that the earlier rule G now becomes redundant). Feature constraints such as that if $V[\text{aux}]$ is the perfective *have*, then XP must be a VP with feature *-en*, are assumed to be enforced through subcategorization features on the $V[\text{aux}]$, such as $V[\text{aux}, \text{VP}[-\text{en}]]$.

while that of $perf_2$ is not. The two operators are introduced to account for certain differences between present and past perfect (not discussed further here), discernible in such contrasts as “*John has left yesterday” *versus* “John had left the day before”; “*John has been sleeping a minute ago” *versus* “John had been sleeping a minute before” (cf., Bennett and Partee 1978). Deindexicalization of $perf_1$ is carried out in accordance with the following rule:²⁵

$$\text{Perf}(1): [(\text{perf}_1 \Phi)_T ** \eta] \leftrightarrow [[[\eta \text{ same-time Ref}'_T] \wedge (\exists e_T: [e_T \text{ until } \eta] [\Phi_{O_1 T} @ e_T])] ** (F \eta)],$$

where Ref'_T denotes the predecessor of Ref_T at the focus of T , $[[F]] = '$, i.e., $(F \eta)$ means “the factual content of situation η ,” and $[\Phi @ e]$ is understood as equivalent to $[\Phi * e]$ (Φ holds throughout e) for *stative* Φ and to $[(\exists e_1: [e_1 \text{ subep } e] [\Phi * e_1]) * (F e)]$ (Φ holds sometime during e) for *nonstative* Φ (as per clause 18 of the formal semantics in the previous section).

We think that this rule, together with a similar one for $perf_2$, solves a number of problems in the interaction of perfectives with tense, adverbials, and the aspectual class of the complement (without resorting to separate methods for the various forms of tensed and untensed perfective, perfect progressive, etc., as is often done in computational linguistics.) One we should go into briefly is that of accounting for the contrast between sentence pairs like

- (21) John has been ill;
- (22) John has become well.

In (21), the episode described by the *have*-complement, John's being ill, extends to the (speaker's) present in the preferred reading.²⁶ In (22), however, the corresponding episode, John's becoming well, *precedes* the present. This is accounted for by the above stipulation that for stative Φ , $[\Phi @ e_T]$ requires that the state described by Φ occupy the entire temporal extent of e_T , whereas for nonstative Φ , the episode described by Φ need only occupy a temporal segment of e_T . Thus, since Perf also specifies that e_T lasts *until* the (new) reference episode η , stative “goings-on” extend all the way to the reference time while nonstative ones need not. Also, Perf , in conjunction with our tense tree mechanism, supplies properly aligned reference times in examples like “Mary looked at

²⁵For those acquainted with Reichenbach's theory of tense, old episode Ref'_T and new episode η both lie at the “reference time”; e_T temporally contains the “event time,” but extends all the way to the reference time. No “speech time” is involved in perfective, *per se*, according to our treatment.

²⁶In the non-preferred reading, John's episode of being ill strictly *precedes* the present. We get such a reading by allowing a “lexical extension rule” to be applied to any lexical verb marked as stative, transforming it into a (less readily available) nonstative verb. The corresponding semantic transformation applies an operator *start-do-stop* to the logical translation, yielding such nonstative predicates as (start-do-stop ill). We think there is independent evidence for availability of (nonpreferred) nonstative readings for all stative verbs, such as the possibility of existential (nongeneric) readings of sentences such as “Vacationers are on the beach,” “Poor people live in those row houses,” or “Criminals own assault rifles.”

the picture; John *had painted it* twelve years ago"; "John thinks that Mary *has left*"; and "John inferred that Mary *had left*."

Conjunctions

Concerning deindexicalization of conjunctions, the most straightforward examples are ones like "John took the money and left," where a tensed verb heads each conjunct. After scoping of the connective (and other operators), the following rule can be applied:

$$\Lambda(1): [\Phi \wedge \Psi]_T \leftrightarrow [\Phi_T \wedge [\Psi_{\Phi \cdot T}]]_T.$$

Roughly speaking, $\Phi \cdot T$ is the tense tree resulting from adding Φ to a discourse context with initial tense tree T . In this way, Ψ is processed relative to a context structure which has already been modified by Φ , as is necessary for making explicit any implicit temporal relations. Thus, the " \cdot " function describes how the tense tree changes. It can be specified recursively with such rules as

$$\begin{aligned} (\text{past } \Phi) \cdot T &= \uparrow(\Phi \cdot (\circ \swarrow T)), \\ (\text{perf}_1 \Phi) \cdot T &= \uparrow(\Phi \cdot (\circ \downarrow T)), \\ (\text{subpast}_1 \Phi) \cdot T &= \Phi \cdot (\circ T), \end{aligned}$$

where $\uparrow T$ is the result of moving the focus of T upward one step. However, details of this need not detain us here. In essence, $(\Phi \cdot T)$ yields the tree one would expect if the shift and store operations specified in the deindexicalization rules were carried out on a *global* data structure (and the focus reset to the node originally focused in T).

A more subtle case of conjunction is the one in which there is just one tensed verb, as in "John lost his job and his wife." Assuming that tense receives widest scope, there is just one tense operator, namely *past*, and this leads to a logical form containing a conjunctive formula characterizing an episode, of the sort $[[\Phi \wedge \Psi] ** \eta]$. The deindexicalization rule in this case is roughly

$$\Lambda(2): [[\Phi \wedge \Psi] ** \eta]_T \leftrightarrow [((\exists e_T: [e_T \text{ subep } \eta] [\Phi_{\circ T} * e_T]) \wedge (\exists e_{T'}: [e_{T'} \text{ subep } \eta] [\Psi_{\circ T'} * e_{T'}]) ** (F \eta))]$$

where $T' = \Phi \cdot (\circ T)$. We have suppressed some slight complications having to do with the fact that e_T and $e_{T'}$ need to be stored in the tense tree as episodes "subordinate" to η , rather than on a par with it. (This can be done with operators $\$$ and \pounds for opening and closing subepisode lists.) The main point we are making here is that conjunction introduces subepisodes, enabling other rules (not shown here) to make any implicit temporal relations between the subepisodes explicit.

Quantifiers

There is a similar phenomenon for sentences involving quantifiers. For instance, in "Everyone congratulated Mary," we can intuitively see that there is not only an "overall" episode covering all the congratulations, but also a set of individual congratulation

episodes, one for each person in question. In "Everyone congratulated Mary, which made her happy," the (preferred) cause of the happiness episode is the "overall" congratulatory episode. But in "Everyone who met Mary immediately congratulated her," the (preferred) reading involves temporal proximity of *individual* meeting events and congratulations. The following rule for monotonic increasing quantifiers (based on semantic rules for quantifiers, 12 and 13) introduces such individual episodes (neglecting certain refinements having to do with surface order):

Quant: For $Q \in \{\forall, \text{Most}, \text{Many}, \text{A-Few}, \dots\}$:

$$[(Q\alpha:\Phi\Psi)_T ** \eta] \leftrightarrow [(Q\alpha:(\exists e:[e \text{ during } \eta][\Phi_T * e]) \\ (\exists e_{T'}:[e_{T'} \text{ subep } \eta] [[\Phi_T \& \Psi_{OT'}] * e_{T'}])]) ** (F \eta)],$$

where $T' = \Phi \cdot T$. Again, $e_{T'}$ should actually be stored as an episode subordinate to η (using \$ and ϕ). On the other hand, the episode e corresponding to (an instance of) the quantifier restriction is not stored in the tense tree at all, according to this rule. This is intended to account for the intuition that NPs do not uniformly provide episodic referents in the way VPs do. (For instance, the referent of *it* is much less apparent in "Most fathers enjoy it" than in "John is a father and enjoys it.") Of course, a noun phrase like *everyone who met Mary* introduces an episodic referent through the relative clause, and this referent can be related to the main-clause episode e_T through tense, aspect, and adverbials.

Adverbials

Our grammar fragment also handles various adverbs and NP/PP-adverbials in combination with tense and aspect and negations. We take adverbials to be uniformly *VP-adverbials* at the level of syntax. (Initial adverbials in sentences like "Yesterday John left" are treated as topicalized.) The following are some of the main rules for such adverbials, with some relevant lexical rules. Note that adverbials are divided into pre-VP and post-VP adverbials, and those that can be either, with feature mod-VP. In the translation of adverbials, we use operators like *adv-e* and *adv-f* that turn a 1-place predicate over episodes into a sentence modifier, *adv-a* that turns a 1-place predicate over actions into a predicate modifier, and *adv-m* that turns a 1-place predicate into a predicate modifier.²⁷ Meaning postulates later apply the predicates which are the arguments of these operators to episodes and actions, respectively.

VP Adjuncts

$VP \rightarrow \text{ADVL}[\text{pre-VP}] VP; (\text{ADVL}' VP')$

$VP \rightarrow VP \text{ADVL}[\text{post-VP}]; (\text{ADVL}' VP')$

Adverbs

²⁷ *adv-e*, *adv-f*, *adv-a*, and *adv-m* stand for episode-modifying adverbial, frequency adverbial, action-modifying adverbial, and manner adverbial, respectively.

$ADVL \rightarrow ADV[\neg\text{not}]; ADV'$
 $ADV[\text{pre-VP}] \rightarrow \text{certainly}; \lambda P \lambda x ((-ly \text{ certain}) [x P])$
 $ADV[\text{mod-VP}] \rightarrow \text{soundly}; (\text{adv-m sound})$
 $ADV[\text{mod-VP}] \rightarrow \text{frequently}; \lambda P \lambda x ((\text{adv-f frequent}) [x P])$

NP/PP Adverbials

$ADVL[\text{post-VP}] \rightarrow NP[\text{time}]; \lambda P \lambda x ((\text{adv-e (during NP')}) [x P])$
 $NP[\text{time}] \rightarrow \text{yesterday}; \text{yesterday}$
 $ADVL[\text{post-VP}] \rightarrow PP[\text{ep-mod}]; \lambda P \lambda x ((\text{adv-e PP'}) [x P])$
 $PP \rightarrow P NP; (P' NP')$
 $P[\text{ep-mod}] \rightarrow \text{for}; \text{lasts}^{28}$

As an illustration of the effect of the rules shown so far, here is the initial translation of a sentence with four adverbials (braces indicate assumed phrase structure):

- (23) {John {certainly {{{slept soundly} for eight hours} yesterday}}}
 $\Rightarrow ((-ly \text{ certain}) ((\text{adv-e (during yesterday)}) ((\text{adv-e (lasts (K1 (eight hour))))$
 $[John ((\text{adv-m sound}) <\text{past sleep}>)])))).$

This is easily verified. The unscoped *pres* operator now needs to be "raised" to some sentential level. There are several possibilities, but the preferred scoping is:

- (24) ((-ly certain)
 (past ((adv-e (during yesterday))
 ((adv-e (lasts (K1 (eight hour))))
 [John ((adv-m sound) sleep)])))).

Next, the following are some additional deindexicalization rules adequate for complete deindexicalization of (24).

Deindexicalization Rules

1. For α a modal operator and Φ a formula:
 $(\alpha \Phi)_T \leftrightarrow (\alpha \Phi_T)$
2. For π, Π non-modal 1-place predicates and Φ a formula:
 $((\text{adv-e } \pi) \Phi)_T \leftrightarrow ((\text{adv-e } \pi_T) \Phi_{\pi.T})$
 $((\text{adv-m } \pi) \Pi)_T \leftrightarrow ((\text{adv-m } \pi_T) \Pi_{\pi.T})$

We also have rules for deindexicalizing indexical terms, e.g., $\text{yesterday}_T = (\text{yesterday-rel-to Now}_T)$, where Now_T is the most recently added episode at the root node of T .

Applying the above rules (together with the *Past*-rule) to (24), we get the following completely deindexicalized formula:

²⁸E.g., (lasts (K1 (one hour))), which means "for one hour."

(25) ((-ly certain)
 (∃e1:[e1 before u1]
 [((adv-e (during (yesterday-rel-to u1)))
 ((adv-e (lasts (K1 (eight hour))))
 [John ((adv-m sound) sleep))]] ** e1))),

where u1 is the utterance of sentence (23).

Finally here are some meaning postulates that could be applied to (25). We implicitly take these to be *necessitated* (i.e., prefixed with \Box).

Meaning Postulates

1. For π a 1-place modal predicate and Φ a formula:
 $((-ly \pi) \Phi) \leftrightarrow [(That \Phi) \pi]$
2. For τ, η terms, π, Π predicates, Φ a formula, and $op \in \{*, **\}$:
 $[((adv-e \pi) \Phi) op \eta] \leftrightarrow [([\eta \pi] \wedge [\Phi * \eta]) op (F \eta)]^{29}$
 $[[\tau ((adv-m \pi) \Pi)] op \eta] \leftrightarrow [([\tau | \eta] (in-manner \pi)) \wedge [[\tau \Pi] * \eta]] op (F \eta)]^{30}$

Applying these MPs to (25), we get the final representation:

(26) [(That (∃e1:[e1 before u1]
 [[e1 during (yesterday-rel-to u1)] ∧
 [[e1 lasts (K1 (eight hour))] ∧
 [([John | e1] (in-manner sound)) ∧
 [John sleep] * e1]] * (F e1))] ** (F e1))]
 certain],

which may be further simplified into³¹

(27) [(That (∃e1:[e1 before u1]
 [[e1 during (yesterday-rel-to u1)] ∧

²⁹E.g., $\pi = \lambda e[e \text{ lasts } (K1 \text{ (two hour)})]$, $\lambda e[e \text{ during yesterday}]$, $\lambda e[e \text{ in Central-Park}]$, ...,
 or, equivalently, $(\text{lasts } (K1 \text{ (two hour)}))$, $(\text{during yesterday})$, (in Central-Park) , ...

Actually, this rule depends on certain "aspectual constraints" being satisfied: *state operators* must be applied only to *stative formulas*. E.g., a *durative* operator like $(adv-e \text{ (lasts } (K1 \text{ (two hour)})))$ is a state operator, and in this rule, cannot be applied to a non-stative formula like $[John \text{ sneeze}]$.

³⁰E.g., $\pi = \text{forceful ("John forcefully expressed his opinion")}$,
 $\text{light ("John kissed Mary lightly on the cheek")}$, ...

Note: *in-manner* is a predicate modifier that turns a predicate over ordinary individuals into a predicate applicable to actions.

³¹The following axiom schemas are used in the simplification process.

For Φ, Ψ atemporal formulas, η a term, and $op \in \{*, **\}$:
 $[[\Phi \wedge \Psi] op \eta] \leftrightarrow [([\Phi * \eta] \wedge [\Psi * \eta]) op \eta]$
 $[[\Phi * \eta] op \eta] \leftrightarrow [\Phi op \eta]$

[[e1 lasts (K1 (eight hour))] \wedge
 [[[John | e1] (in-manner sound)] \wedge
 [[John sleep] * e1]]]] ** (F e1)))]
 certain].

5 Making Inferences with Episodic Logic

Our inference rules fall into three broad categories, namely, (i) basic inference rules including rule instantiation and its dual goal reduction, (ii) narrative inference rules such as causal connection, temporal succession and state persistence, and (iii) simulative inference rules. All of the rules allow for epistemic probability bounds on premises and for use of generic conditionals.

We have already briefly described Rule Instantiation (RI) which is heavily used in input-driven inference; its dual Goal Chaining (GC) similarly dominates goal-driven inference. We have indicated that RI has *modus ponens* as special case, but also allows instantiation of generic conditionals. In fact, it allows arbitrarily many "minor premises" to be matched against arbitrarily deeply embedded subformulas of a rule. (Apart from its avoidance of skolemization, it resembles Andrews' *general matings* (1981) and Bibel's *connections* (1979).) Schematically, the rule is

RI (Rule Instantiation)

– Non-probabilistic version –

For $R(\Phi_1, \dots, \Phi_m, \Phi'_1, \dots, \Phi'_n)$, $\Psi_1, \dots, \Psi_m, \neg\Psi'_1, \dots, \neg\Psi'_n$, formulas with bound variables standardized apart, and with all Φ_i 's occurring negatively in $R(\Phi_1, \dots, \Phi_m, \Phi'_1, \dots, \Phi'_n)$, and all Φ'_i 's occurring positively in it:

$$\frac{R(\Phi_1, \dots, \Phi_m, \Phi'_1, \dots, \Phi'_n) \quad \Psi_1, \dots, \Psi_m, \neg\Psi'_1, \dots, \neg\Psi'_n}{R_\sigma(\top, \dots, \top, \perp, \dots, \perp)}$$

where substitution σ unifies the Φ_i with corresponding Ψ_i and Φ'_i with corresponding $\neg\Psi'_i$. $R_\sigma(\top, \dots, \top, \perp, \dots, \perp)$ is then simplified to eliminate the truth values \top and \perp .

The substitution σ applies to certain matchable variables which are \forall -quantified by a positively occurring quantifier, or \exists -quantified by a negatively occurring quantifier, in $R(\Phi_1, \dots, \Phi_m, \Phi'_1, \dots, \Phi'_n)$ or one of the Ψ_i . A subformula occurs positively if it lies within an even number of negations, where " \neg ", conditional antecedents, and \forall -quantifier restrictions count as negation, and similarly, for "negatively occurring." Computing $R_\sigma(\top, \dots, \top, \perp, \dots, \perp)$ involves elimination of quantifiers of variables replaced by σ , e.g., if b is substituted for x , then $(\forall x:\Phi\Psi)_{b/x}$ becomes $[\Phi_{b/x} \rightarrow \Psi_{b/x}]$. A probabilistic

version of RI results when the Ψ_i or $\neg\Psi'_i$ are allowed to have non-unit lower epistemic probabilities and/or R is a generic conditional.

A rule instantiation typically instantiates the complete antecedent of a rule, and infers the particularized consequent. However, it may only match part of the antecedent, or match part or all of the consequent (giving a “contrapositive” inference). For example, consider rule

$$(\forall x[[x \text{ whale}] \rightarrow \neg[x \text{ fish}]]),$$

i.e., no whale is a fish. Here x is a matchable variable since it is quantified by a positively occurring \forall -quantifier. $[x \text{ whale}]$ and $[x \text{ fish}]$, called, say, Φ_1 and Φ_2 , occur negatively, while their negations occur positively. Suppose now we have the assertion [Moby-Dick whale] in the knowledge base. Then, substitution Moby-Dick/ x unifies [Moby-Dick whale] with Φ_1 , with result

$$\top \rightarrow \neg[\text{Moby-Dick fish}],$$

which is simplified to $\neg[\text{Moby-Dick fish}]$. Conversely, if we have [Wanda fish] in the knowledge base, unification with Φ_2 yields the inference $[\text{Wanda whale}] \rightarrow \neg\top$, i.e., $\neg[\text{Wanda whale}]$.

As another example, consider the rule

$$(\forall x[[x \text{ large-carnivore}] \rightarrow ((\exists e[[x \text{ hungry}] ** e]) \rightarrow (\exists e'[[e' \text{ sametime } e] \wedge [[x \text{ dangerous}] ** e']]]))],$$

i.e., Every large carnivore is dangerous whenever it is hungry. Here x and e are matchable, x being quantified by a positively occurring \forall -quantifier, and e by a negatively occurring \exists -quantifier. Also, note that $[x \text{ large-carnivore}]$ and $[[x \text{ hungry}] ** e]$, called, say, Φ_1 and Φ_2 , occur negatively in this rule while $[e \text{ sametime } e']$ and $[[x \text{ dangerous}] ** e']$ occur positively. Suppose now we get input [C11 wolf] and $[[\text{C11 hungry}] ** E7]$. Using additional facts from knowledge base, in this case, hierarchical knowledge implicit in (type) specialists, we can infer [C11 large-carnivore] from [C11 wolf]. Then, substitution $\{\text{C11}/x, E7/e\}$ unifies [C11 large-carnivore] and $[[\text{C11 hungry}] ** E7]$ with Φ_1 and Φ_2 , with result

$$\top \rightarrow [\top \rightarrow (\exists e'[[e' \text{ sametime } E7] \wedge [[\text{C11 dangerous}] ** e']]]],$$

which is simplified to $(\exists e'[[e' \text{ sametime } E7] \wedge [[\text{C11 dangerous}] ** e']])$. This process amounts to making the inference “The wolf is dangerous at the time of E7 (the episode of being hungry).”

The generalization of RI we are using for probabilistic inference is essentially the same as the nonprobabilistic version, with the antecedent of a generic conditional counting as a negative environment, and the consequent as positive. Additional rules of simplification are needed, including: $\top^q \rightarrow_p \Phi$ becomes Φ^{pq} , $\perp^q \rightarrow_p \Phi$ becomes \top , $\Phi \rightarrow_p \top^q$ becomes \top , and $\Phi \rightarrow_p \perp^q$ becomes $(\neg\Phi)^{pq}$. (These rules are justifiable in terms of probability theory for $q = 1$ or 0 , but for intermediate q they merely reflect our intuitions

at this point.) As an illustration, if the two conditionals in the "large-carnivore" example above had weights p and q , the previous conclusion would be obtained with lower probability pq .

In practice, RI is implemented roughly as follows. A newly inferred conclusion, corresponding to one of the Ψ_i or $\neg\Psi'_i$, is used to index to the rule R . An initial determination is then made whether the instantiation is likely to succeed and yield a useful result. If the decision is to instantiate, then the attempt to do so is performed by a recursive algorithm applied to R , which actively seeks to find appropriate Ψ_i and $\neg\Psi'_i$ instances in the knowledge base to unify with negatively and positively occurring subformulas of R . Actually, as was indicated in the second illustration, the Ψ_i and $\neg\Psi'_i$ need not even occur explicitly in the knowledge base. They may be inferred by specialists for type taxonomies, temporal relations, or other special classes of relations, or by a limited amount of Prolog-like backchaining. Before turning to goal-directed inference rules, we should mention λ -conversion and substitution of equals for equals as further deductive rules available in our logic.

For goal-directed inference (e.g., in response to questions) two general methods are available. The first, **Goal Chaining (GC)**, is the dual of RI. For comprehensibility, we state only a special (but frequently encountered) case.

GC (Goal Chaining)

For $R(\Phi)$, Ψ standardized formulas where Φ is a positively occurring subformula of $R(\Phi)$:

$$\frac{R(\Phi), \Psi}{\neg R_{\bar{\sigma}}(\perp)}$$

where $\bar{\sigma}$ differs from σ (in RI) in that it treats variables of Ψ with positively occurring \exists -quantifiers or negatively occurring \forall -quantifiers as matchable.

Like RI, this is a very general chaining rule, allowing not only chaining from rule consequents to antecedents, but from *any* positively occurring subformula to the rest of $R(\Phi)$ (negated and suitably instantiated).

For example, consider a rule R ,

$$(\exists x[x \text{ wolf}]) \rightarrow [[x \text{ meat-eater}] \wedge [x \text{ forest-dweller}]],$$

and a goal Ψ ,

$$(\exists x_1[x_1 \text{ forest-dweller}]).$$

In R , subformulas $[x \text{ meat-eater}]$ and $[x \text{ forest-dweller}]$ occur positively; in Ψ , x_1 is matchable as it is quantified by a positively occurring \exists -quantifier. Thus, via substitution x_1/x , we can unify $[x_1 \text{ forest-dweller}]$ with $[x \text{ forest-dweller}]$, and get

$$\neg((\exists x[x \text{ wolf}]) \rightarrow [[x \text{ meat-eater}] \wedge \perp]),$$

which is simplified to $\neg((\exists x[x \text{ wolf}]) \rightarrow \perp)$, i.e., $(\exists x[x \text{ wolf}])$. This process amounts to reducing the question "Is there a forest-dweller?" to "Is there a wolf?," using knowledge "A wolf is a meat eater and a forest dweller."

As another example, consider rule R ,

$$(\exists x[x \text{ wolf}]) \rightarrow [[x \text{ gray}] \vee [x \text{ black}]],$$

and a goal Ψ ,

$$(\exists x_1[x_1 \text{ gray}]).$$

In R , subformulas $[x \text{ gray}]$ and $[x \text{ black}]$ occur positively; in Ψ , x_1 is matchable much as before. Thus, via substitution x_1/x , we can unify $[x_1 \text{ gray}]$ with $[x \text{ gray}]$, and get

$$\neg((\exists x[x \text{ wolf}]) \rightarrow [[x \text{ black}] \vee \perp]),$$

which is equivalent to $\neg((\exists x[x \text{ wolf}]) \rightarrow [x \text{ black}])$, i.e., $(\exists x[[x \text{ wolf}] \wedge \neg[x \text{ black}]])$. This process amounts to reducing the question "Is there a gray thing?" to "Is there a wolf that is not black?," using knowledge "A wolf is either gray or black."

The general version of GC, like the general RI rule, allows arbitrarily many subsidiary knowledge base facts to be invoked in the process of chaining from the given goal to a subgoal. For instance, in the second example above, knowledge base fact $[C11 \text{ wolf}]$, $\neg[C11 \text{ black}]$ would have led to immediate success (i.e., subgoal \top). Also, probabilities are handled much as in RI. (A subgoal Φ^p is interpreted as meaning that if Φ can be proved with probability q , then the original goal is established with probability pq .)

The second class of goal-directed methods (not yet implemented) consists of standard natural deduction rules such as proving a conditional by assuming the antecedent and deriving the consequent; or proving a negative formula by assuming the positive and deriving a contradiction; or proving a universal by proving an "arbitrary instance" of it. Such rules are needed for completeness, since goal chaining cannot prove valid formulas such as $\Phi \rightarrow \Phi$. An interesting future possibility, in the case of proofs involving assumption-making, is to activate input-driven inferencing (primarily, RI) once an assumption has been made, so that its important consequences will be worked out, making it easier to complete the goal-directed proof.

In our implementation, we use a sophisticated agenda-driven control structure for goal chaining (largely borrowed from ECONET - see de Haan & Schubert 1986) with goals ranked according to estimated difficulty and new knowledge accessed via concept and topic hierarchies. The aim here is not so much theorem proving power *per se*, but the ability to get at the *relevant* knowledge in a large knowledge base. Despite the lack of natural deduction rules, all the examples which ran in the original resolution-based system are handled by ECOLOGIC (in addition to new examples involving **, λ -abstraction, and so on).

A remaining problem is, of course, the principled handling of probabilities. The state of the art in probabilistic inference (e.g., Pearl 1988, Bacchus 1988a) is not such as to

provide concrete technical tools for a logic as general as episodic logic. We are, however, successfully using a “noncircularity principle” which prevents the same knowledge from being used twice to “boost” the probability of a particular conclusion. This is done by keeping track of the support set in a probabilistic inference process. Apart from this, we use independence assumptions where there are no known dependencies, and manipulate lower probabilities in accord with the laws of probability.

Finally, we mention some narrative and simulative inference rules (yet to be implemented).

CC (Causal Connection)

$$\frac{[\eta' \text{ successor-ep } \eta], [\eta \text{ event}], [\eta' \text{ event}], \neg(\exists x[[x \mid \eta'] \text{ volitional-action}])}{[\eta \text{ cause-of } \eta']^6}$$

For example, given a fragment “John greeted Mary; Mary was startled,” we conclude that John’s greeting is the cause of Mary’s being startled, with minimal degree of belief .6. (Note that since the conclusion is probabilistically qualified, it can be overridden by other considerations, such as plausibility or other “coherence” relations. For instance, in “The sun set; the moon rose,” we would want the causal inference to be overridden.) We should remark that we are in the process of reformulating narrative inference rules as generic conditionals to be used in the same way as narrative domain knowledge. For example, the above rule can be reformulated as a generic conditional with η and η' replaced by existentially quantified variables, and the generic conditional carrying weight .6. (More accurately, all of the premises and the conclusion should be embedded within *performative* predications, such as “text-source asserts that ...”.)

An example of a simulative inference rule is:

SIM (Simulative Reasoning)

For τ , an individual term; $\pi \in \{\text{learn, discover, remember, find, realize ...}\}$; Φ and Ψ , formulas; η , an episodic term; \mathcal{K} , a set of facts which are “common knowledge” (shared by the reasoner and τ):

$$\frac{\begin{array}{c} [[\tau \pi (\text{That } \Phi)] ** \eta] \\ \Phi \vdash_{\mathcal{K}} \Psi \end{array}}{(\exists e: [[e \text{ right-after } \eta] \wedge [\eta \text{ cause-of } e]])} \\ [[\tau \text{ infer } (\text{That } \Psi)] ** e]$$

where $\vdash_{\mathcal{K}}$ means “follows automatically by input-driven inference, with use of (only) the additional premises \mathcal{K} in the knowledge base.”

Note that in this *simulative reasoning* rule, the question answerer’s *own* ability to infer Ψ from Φ via common knowledge is being attributed to τ , the agent in the antecedent.

6 Episodic Logic and Narrative Understanding

In outline, story understanding on our view involves the following interleaved steps for each new sentence: (i) parsing and initial logical translation; (ii) disambiguation (including quantifier scoping and anaphora resolution); (iii) deindexicalization; (iv) application of all three types of inference rules to the deindexicalized-translated input, in combination with stored knowledge (meaning postulates, generic conditionals, and other general and specific knowledge); among other things, this may generate new predictions and explanations; (v) matching of previous predictions and explanations with new ones.

We have put these ideas to the test in two ways: first, by hand-simulating the inference process for a small fragment of the story of *Little Red Riding Hood*. We also have a prototype implementation which accepts logical-form inputs and performs many of the inferences we have alluded to and is able to answer simple questions (Schubert *et al.* 1989). In this section, we will show how our logic would allow the system to account for the wolf's decision not to eat Little Red Riding Hood right away when he first met her, given a brief excerpt from the story as follows.

*In the forest, Little Red Riding Hood met a wolf.
The wolf would have very much liked to eat her, but
he dared not do so on account of some woodcutters nearby.*

Processing this fragment requires extensive reasoning including inferences based on meaning postulates, predictive inferences, explanatory inferences and simulative inferences. For example, to understand the third sentence, one should be able to explain why the wolf decided against eating Little Red Riding Hood, and how the presence of woodcutters nearby affected the wolf's decision. So, one has to know that when some agent *dares* not do something, he must think it possible that his attempt to do it would result in something unpleasant to himself; then one has to simulate his reasoning process to guess what unpleasant consequences he anticipates.

Depending on the degree of sophistication of the knowledge possessed, people may explain the wolf's decision in various ways. Correspondingly, depending on the kind of knowledge provided, our inference machinery can produce various lines of reasoning; this includes the following, relatively simple line of reasoning

- Attacking a child is extremely wicked.
- Trying to eat a living creature involves attacking it, and such an attack is conspicuous and likely to be noticed by nearby people.
- Doing something extremely wicked is likely to bring severe punishment, if noticed by anyone.
- So, if the wolf tries to eat Little Red Riding Hood, the nearby woodcutters may notice it, and he is likely to be severely punished for it.

Or, the more sophisticated version

- When a predatory animal eats a non-predatory creature of comparable size while the creature is conscious, the predator attacks it as a preparation for eating it.
- The wolf would attack Little Red Riding Hood before eating her.
- Attacking a person is a conspicuous action, and is likely to be noticed by nearby people.
- If people notice a predatory animal attacking a person, they will most probably want to rescue the person from the animal.
- To rescue a person from a predatory animal, one may kill it.
- Thus, the woodcutters may kill the wolf.

Upon reaching the conclusion that it is possible that the wolf might be killed or severely punished, our inference machinery attributes its own ability to infer that conclusion to the wolf (this is due to our rules of simulative reasoning). Then it is easily explained why the wolf decided against eating Little Red Riding Hood right then and there.

In the following we show in detail that part of the reasoning process reaching the conclusion "The wolf may be severely punished." The inferences we show have actually been generated by the current implementation (in nearly the same order). The control structure is designed to systematically combine each new clause with relevant meaning postulates and other general knowledge. All of the inferences are based on the explicit, formalized rules of inference we introduced earlier. (Simulative and narrative inferences are not yet made, but are not needed here.) After listing meaning postulates and world knowledge, we show the logical translation of the story and the reasoning process.

Meaning Postulates

M1. To walk, to attack someone, to try to do something, to die, etc., are types of actions.

For Π an action predicate:

$\square [(K_a \Pi) \text{ action-type}]$

An "action predicate" is an expression $(\pi \tau_1, \dots, \tau_{n-1})$, where π is an n -adic atomic action predicate, $n \geq 1$, and $\tau_1, \dots, \tau_{n-1}$ are terms.

M2 (A meaning postulate regarding actions).

For Π an action predicate:

$\square (\forall x(\forall e[[x \Pi] ** e] \leftrightarrow [[x | e] \text{ instance-of } (K_a \Pi)]))$

For example,

$[[\text{John eat}] ** E1] \leftrightarrow [[\text{John} | E1] \text{ instance-of } (K_a \text{ eat})].$

Note that $[\text{John} | E1]$ is an action, not just an arbitrary individual-episode pair, so that John is the agent of that action.

- M3. If there is a collection of things of some type, then there is a thing of that type which belongs to that collection (we regard collections as non-empty by definition).

For Π a monadic predicate:

$$\Box (\forall x:[x \text{ (coll } \Pi)] (\exists y:[y \text{ in } x][y \Pi]))$$

"coll" is a function that maps a predicate applicable to things into a predicate applicable to collections of things.

World Knowledge

- K1. For a creature to attack a child is extremely wicked.

$$(\exists x:[x \text{ creature}](\exists y:[y \text{ child}](\exists e[[x \text{ attack } y] ** e]])) \\ \rightarrow .9 [[x | e] ((\neg \text{ly extreme}) \text{ wicked})]$$

- K2. Trying to eat any living creature involves attacking it.

$$(\forall x:[x \text{ alive}] \wedge [x \text{ creature}]) \\ [(K_a (\text{try } (K_a (\text{eat } x)))) \text{ involve } (K_a (\text{attack } x)))]$$

- K3. If one type of action involves another, then any creature doing an instance of the first will do an instance of the second during it.

$$(\forall a1:[a1 \text{ action-type}] \\ (\forall a2:[a2 \text{ action-type}] \wedge [a1 \text{ involve } a2]) \\ (\forall x:[x \text{ creature}] \\ (\forall e1:[x | e1] \text{ instance-of } a1) \\ (\exists e2:[e2 \text{ during } e1] \\ [[x | e2] \text{ instance-of } a2]]))))$$

- K4. For a sizable creature to attack a sizable thing is conspicuous (relative to a human observer).

$$(\exists x:[x \text{ person}] \\ (\exists y:[y \text{ creature}] \wedge \neg [y \text{ tiny-rel-to } x]) \\ (\exists z:[z \text{ creature}] \wedge \neg [z \text{ tiny-rel-to } y]) \\ (\exists e[[y \text{ attack } z] ** e]])) \\ \rightarrow .9 [[y | e] \text{ conspicuous-to } x]$$

By contrast, for an ant to attack something would not be conspicuous to a human.

- K5. If a creature performs a conspicuous action within plain sight of a person, that person is likely to notice that action.

$(\exists x:[x \text{ creature}]$
 $(\exists y:[y \text{ person}]$
 $(\exists e1:[x \text{ within-plain-sight-of } y] ** e1]$
 $(\exists e2:[e2 \text{ during } e1] [[x | e2] \text{ conspicuous-to } y]]))$
 $\rightarrow .6 (\exists e3:[e3 \text{ during } e2] [[y \text{ notice } [x | e2]] ** e3])$

- K6. Doing something extremely wicked may bring severe punishment from some group of people, if noticed by anyone.

$(\exists x:[x \text{ creature}]$
 $(\exists e1:[x | e1] ((-ly \text{ extreme}) \text{ wicked}))$
 $(\exists y:[y \text{ person}]$
 $(\exists e2 [[y \text{ notice } [x | e1]] ** e2]]))$
 $\rightarrow .3 (\exists z:[z \text{ (coll person)}]$
 $(\exists e3:[e2 \text{ cause-of } e3]$
 $[[z ((-ly \text{ severe}) (\text{punish } x))] ** e3]))$

- K7. A human is not tiny relative to a wolf, and vice versa.

$(\forall x:[x \text{ human}] (\forall y:[y \text{ wolf}]$
 $[\neg[x \text{ tiny-rel-to } y] \wedge \neg[y \text{ tiny-rel-to } x]]))$

- K8. If a creature is near a person and not tiny relative to the person, it is probably within plain sight of the person. (This could be improved by assuming that we are dealing with a *daytime* episode in an open setting.)

$(\exists x:[x \text{ person}]$
 $(\exists y:[y \text{ creature}] \wedge \neg[y \text{ tiny-rel-to } x])$
 $(\exists e1 [[y \text{ near } x] ** e1]))$
 $\rightarrow .6 (\exists e2:[e2 \text{ same-time } e1]$
 $[[y \text{ within-plain-sight-of } x] ** e2])$

- K9. Woodcutters are humans.

$(\forall x:[x \text{ woodcutter}] [x \text{ human}])$

Story

Now, let's work out the possible consequences if the wolf tries to eat Little Red Riding Hood. (We then attribute this reasoning to the wolf.) The relevant assumptions and story facts are as follows (where we use the convention of having variables in lower case, and constants in upper case):

The wolf tries to eat Little Red Riding Hood.

$(\exists e1:[now^{32} \text{ during } e1] (\text{The } x1:[x1 \text{ wolf}][[x1 \text{ try } (K_a (\text{eat LRRH}))] ** e1]))$

By skolemization $\{E1/e1\}$ and reference determination $\{W/x1\}$:

- S1. $[now \text{ during } E1]$
- S2. $[W \text{ wolf}]$
- S3. $[[W \text{ try } (K_a (\text{eat LRRH}))] ** E1]$

Little Red Riding Hood is a girl and alive.

- S4. $[LRRH \text{ girl}]$
- S5. $[LRRH \text{ alive}]$

There are woodcutters nearby.

$(\exists y1:[y1 \text{ (coll woodcutter)}]$
 $(\forall x:[x \text{ in } y1] (\exists e2:[E1 \text{ during } e2] [[W \text{ near } x] ** e2])))$

By Skolemizing $\{C1/y1\}$:

- S6. $[C1 \text{ (coll woodcutter)}]$
- S7. $(\forall x:[x \text{ in } C1] (\exists e2:[E1 \text{ during } e2] [[W \text{ near } x] ** e2]))$

Assume the following type-hierarchical knowledge is available (at least indirectly, via a type "specialist"):

- S8. $[W \text{ creature}]$
- S9. $[LRRH \text{ child}]$
- S10. $[LRRH \text{ human}]$
- S11. $[LRRH \text{ creature}]$

Reasoning Process

- Note that simple time inferences such as
 $[E1 \text{ during } E2] \wedge [E2 \text{ during } E3] \wedge [E3 \text{ same-time } E4] \vdash [E1 \text{ during } E4]$
will be taken for granted during the inference process.

³² *now* will be replaced by a term with fixed reference to the time of speech.

- In the following,

RI [A; B] {Subst C/v; Imm-Skol C'/v'}

indicates that the subsequent inference(s) has been made via **Rule Instantiation** of rule B by premise(s) A, with variable substitution C/v, and an existential-variable v' in the inferred formula has been immediately skolemized as C'.

RI [S5, S11; K2] {Subst LRRH/x}:

1. [(K_a (try (K_a (eat LRRH)))) involve (K_a (attack LRRH))]
"Trying to eat LRRH involves attacking her."

RI [; M1] {Subst (try (K_a (eat LRRH)))/Π}:

2. [(K_a (try (K_a (eat LRRH)))) action-type]
"Trying to eat LRRH is an action type."

RI [; M1] {Subst (attack LRRH)/Π}:

3. [(K_a (attack LRRH)) action-type]
"Attacking LRRH is an action type."

RI [S3; M2] {Subst W/x, E1/e, (try (K_a (eat LRRH)))/Π}:

4. [[W | E1] instance-of (K_a (try (K_a (eat LRRH))))]
"The wolf's trying to eat LRRH is an instance of someone's trying to eat LRRH."

RI [2, 3, 1, S8, 4; K3]

{Subst (K_a (try (K_a (eat LRRH))))/a1, (K_a (attack LRRH))/a2, W/x, E1/e1;
 Imm-Skol E2/e2}:

5. [E2 during E1]
6. [[W | E2] instance-of (K_a (attack LRRH))]

RI [6; M2] {Subst W/x, E2/e, (attack LRRH)/Π}:

7. [[W attack LRRH] ** E2]
*"The wolf attacks LRRH."*³³

³³Inferences 5 and 7 could be obtained in one step rather than six by using the following knowledge K2' instead of K2.

K2'. When a creature tries to eat a creature that is alive, he attacks it during that episode.

(∃x:[x creature]
 (∃y:[y alive] ∧ [y creature])
 (∃e1[[x try (K_a (eat y))] ** e1]))
 → (∃e2:[e2 during e1] [x attack y] ** e2))

Specifically,

RI [S8, S5, S11, S3; K2'] {Subst W/x, LRRH/y, E1/e1; Imm-Skol E2/e2}:

5. [E2 during E1]
7. [[W attack LRRH] ** E2]

RI [S8, S9, 7; K1] {Subst W/x, LRRH/y, E2/e}:

8. $[[W \mid E2] ((-ly \text{ extreme}) \text{ wicked})]$ ⁹

Up to here:

The wolf attacks LRRH, and that's extremely wicked.

RI [S10, S2; K7] {Subst LRRH/x, W/y}:

9. $\neg[LRRH \text{ tiny-rel-to } W]$

"LRRH is not tiny relative to the wolf."

RI [S6; M3] {Subst C1/x, woodcutter/II; Imm-Skol C2/y}:

10. [C2 in C1]

11. [C2 woodcutter]

"There is a woodcutter."

RI [11; K9] {Subst C2/x}:

12. [C2 human]

"The woodcutter is a human."

With type-hierarchical knowledge, we get from 12:

13. [C2 person]

RI [12, S2; K7] {Subst C2/x, W/y}:

14. $\neg[W \text{ tiny-rel-to } C2]$

"The wolf is not tiny relative to the woodcutter."

RI [13, S8, 14, S11, 9, 7; K4] {Subst C2/x, W/y, LRRH/z, E2/e}:

15. $[[W \mid E2] \text{ conspicuous-to } C2]$ ⁹

Up to here:

The wolf's attack is conspicuous to the woodcutter.

RI [10; S7] {Subst C2/x; Imm-Skol E3/e}:

16. [E1 during E3]

17. $[[W \text{ near } C2] ** E3]$

"The wolf is near the woodcutter (when he tries to eat LRRH)."

RI [13, S8, 14, 17; K8] {Subst C2/x, W/y, E3/e1; Imm-Skol E4/e2}:

However, our aim is to obtain the desired inferences in narrative understanding from *any* reasonable, intuitively natural way of formulating the relevant world knowledge. K2 is probably more natural than K2', and more importantly, was written down *prior* to detailed consideration of the reasoning process it was intended to support. If we are going to have a robust system whose knowledge base and range of understanding is easily expanded, we cannot afford to "tailor" the syntactic form of the axioms to the inference chains we choose as examples.

18. [E4 same-time E3]
 19. [[W within-plain-sight-of C2] ** E4]⁶
"The wolf is likely to be within plain sight of the woodcutter."

RI [S8, 13, 19, (5, 10, 18), 15; K5] {Subst W/x, C2/y, E4/e1, E2/e2;
 Imm-Skol E5/e3}:

20. [E5 during E2]
 21. [[C2 notice [W | E2]] ** E5]³²⁴
 # Up to here:
 The woodcutter may notice the wolf's attacking LRRH.

RI [S8, 8, 13, 21; K6] {Subst W/x, E2/e1, C2/y, E5/e2; Imm-Skol C3/z, E6/e3}:

22. [C3 (coll person)]
 23. [E5 cause-of E6]
 24. [[C3 ((-ly severe) (punish W))]] ** E6]⁰⁸⁷
 # The wolf may be severely punished by some group of people.

This inference chain can be extended to provide an explanation for the wolf's decision *not* to try to eat LRRH at that point in the story. First, rule K3 would be slightly augmented so as to express the fact that if one action involves another, and that other action has certain consequences, then these are also consequences of the first action. Rule K5 would be similarly augmented to make the "noticing episode" e3 a causal consequence of the episode e2 (or action [x | e2]) noticed. *The "punishing episode," E6, in conclusion 24 would then be inferred to be a consequence of the wolf's attempt to eat LRRH.* Given that being severely punished is very bad, and that agents generally refrain from actions that they think may have very bad consequences for them, we would have an explanation for the wolf's restraint. Note, however, that this requires application of the simulative inference rule (SIM), i.e., we must attribute the above inference chain to the wolf, and draw further conclusions from this attribution.

7 Concluding Remarks

Our logic is probably the most expressive yet brought to bear on the problem of narrative understanding. It makes implicit time and situation dependencies explicit through the use of episodic variables, and admits unbound "anaphoric" variables and the representation of generic conditionals. The expressive power of our logic exceeds that of most natural language systems (*cf.*, Alshawi and van Eijck (1989)).

Our use of episodic variables owes a debt to Davidson (1969), but we can "attach" an episodic variable to any formula, whereas Davidson's method can introduce episodes only for atomic formulas. Thus, for Davidson, there can be no episodes involving *quantification*, such as an episode of everyone in the room looking at Mary, and no episodes

involving *negation*, such as an episode of John not eating anything for ten hours. Yet such episodes can perfectly well be cited as causal antecedents (e.g., sentences (8) and (9) in section 2; or "John did not eat anything for ten hours, and as a result, he was famished"), anaphorically referred to, quantified over, etc.

We should also mention the Situation Calculus of McCarthy and Hayes (1969), whose notion of a situation corresponds exactly to our notion of a (possible) moment of time. As well, the Event Calculus of Kowalski (1986) treats events as individuals as we do, the *occasions* of Firby and McDermott (1987) correspond to our "*", and the "'" operator of Hobbs *et al.* (1986) is similar to our "*."³⁴ Like Davidson, however, all of them handle only atomic formulas and are unable to deal with events involving quantification and logical compounds.

We have provided evidence that our episode-based logical form can provide a clean foundation for story understanding. The main advantages of our approach are as follows:

- (a) The representation of phrase structure is modular and transparent, as is the mapping from phrase structure to episodic logic. The mapping handles many combinations of tense, aspect and adverbials.
- (b) Episodic logic is expressively rich — it allows the content of most English sentences and most world knowledge to be represented in an intuitively comprehensible and formally analyzable manner. Restricted quantifiers, modal operators, nominalization operators, episodic variables, anaphoric variables, and generic conditionals are brought together for the first time in a logic for narrative understanding.
- (c) Being probabilistic, our rules of inference allow evidence for explanations or predictions to be weighed, much as in expert systems.
- (d) All types of linguistic and domain knowledge are strictly separated from parsing and inference control structure, allowing the former to be expanded and revised independently of the latter.
- (e) Hand-simulation of the processing of actual story fragments, and question-answering, indicates that our logical framework is epistemologically adequate for story understanding.

This last claim, about epistemological adequacy, may come as something of a surprise. Whatever happened to scripts, plans, TAUs, TOPs, MOPs, etc.? Are these "higher-level" knowledge structures not essential to story comprehension? We do not doubt that they are. However, we see no sharp divisions between any of them. The more focused the successive stages of a script are on an ultimate goal, the more it resembles a plan. The more abstract its level of description, the more it resembles a TAU or a TOP, and so on. Furthermore, we see no particular obstacle to encoding all of them as axiomatic knowledge in episodic logic, in the manner of the examples in section 2.

³⁴The relationship between their "'" operator and our "*" appears to be $\pi' =_{df} \lambda e \lambda x [[x \pi] * e]$; however, Hobbs *et al.* provide no formal semantics for "'".

For example, the M-BORROW MOP (Dyer 1983: 207) can be cast as a set of generic conditionals along the following lines. If some person x wants to have some object y temporarily, which he knows to be in the possession of some person z , he may well ask z to lend him y and this may induce z to do so, fulfilling x 's goal. If some person x has some object y on loan from some person z , then x is obligated to return y to z , and z will probably want him to do so; etc. We consider the taxonomy of scripts, plans, MOPs, etc., and their elaborate subcategorization, more of a potential guide to control structure – what knowledge is likely to be useful when – than a guide to representation.

Much work remains to be done on our logic – for instance, on the formal semantics of nominalization and propositional attitudes, and of probabilistic inference. As well, we need to compose many more rules of translation and compile a substantial body of knowledge for particular stories. However, our implementation to date has proved to be very gratifying (ECOLOGIC: Schubert *et al.* 1989), and we have incorporated several techniques into ECOLOGIC which were developed for an earlier system based on ordinary first-order logic (ECOSYSTEM: de Haan & Schubert 1986, Miller *et al.* 1987, Miller & Schubert 1988), facilitating efficient deduction (both general and specialized) and fast, selective access to knowledge relevant to a particular set of concepts and topics. “Natural” goal reduction remains to be implemented, but nonetheless the types of questions handled by ECONET are also handled by ECOLOGIC (e.g., “Did anyone have some cake?” or “Does grandmother live in a shoe?”)

Acknowledgements

The authors are indebted to Philip Harrison for many detailed criticisms and suggestions. They are also grateful to Bob Wilensky for some insightful comments on the individuation of events that prompted revisions in the semantics of “**.” Stephanie Miller has been the mainstay of the program development effort for ECOLOGIC. As well, James Allen provided helpful comments, and the University of Alberta Logical Grammar Study Group provided a forum for general discussion of the topic. This research was supported by NSERC Operating Grant A8818 (LKS), an Izaak W. Killam Memorial Scholarship (CHH), the Boeing Co. under Purchase Contracts W-278258 and W-288104, and ONR/DARPA research contract no. N00014-82-K-0193.

References

- [Allen, 1987] J. Allen, *Natural Language Understanding*, Benjamin/Cummings Publ. Co., Reading, Mass., 1987.
- [Alshaw and van Eijck, 1989] H. Alshaw and J. van Eijck, “Logical forms in the Core Language Engine,” In *Proc. of the 27th Annual Meeting of the ACL*, pages 25–32, Vancouver, B.C., Canada, June 26–29, 1989.

- [Andrews, 1981] P. B. Andrews, "Theorem proving via general matings," *JACM*, 28(2):193-214, 1981.
- [Bacchus, 1988a] F. Bacchus, *Representing and Reasoning with Probabilistic Knowledge*, PhD thesis, U. of Alberta, Edmonton, Alberta, 1988.
- [Bacchus, 1988b] F. Bacchus, "Statistically founded degrees of belief," In *Proc. of the 7th Bienn. Conf. of the Can. Soc. for Computational Stud. of Intelligence (CSCSI '88)*, pages 59-66, Edmonton, Alberta, June 6-10, 1988.
- [Barwise, 1989] J. Barwise, *The Situation in Logic*, CSLI: Center for the Study of Language and Information, Stanford, CA, 1989.
- [Barwise and Cooper, 1981] J. Barwise and R. Cooper, "Generalized quantifiers and natural language," *Ling. and Phil.*, 4(2):159-219, 1981.
- [Barwise and Perry, 1983] J. Barwise and J. Perry, *Situations and Attitudes*, MIT Press (Bradford Book), Cambridge, Mass., 1983.
- [Bennett and Partee, 1978] M. Bennett and B. Partee, *Toward the Logic of Tense and Aspect in English*, Indiana U. Linguistics Club, 1978.
- [Bibel, 1979] W. Bibel, "Tautology testing with a generalized matrix reduction method," *Theor. Comput. Sci.*, 8:31-44, 1979.
- [Chierchia and Turner, 1988] G. Chierchia and R. Turner, "Semantics and property theory," *Ling. and Phil.*, 11:261-302, 1988.
- [Davidson, 1969] D. Davidson, "The individuation of events," In N. Rescher et al., editor, *Essays in Honor of Carl G. Hempel*, pages 216-234. Reidel, Dordrecht, Holland, 1969.
- [de Haan and Schubert, 1986] J. de Haan and L. K. Schubert, "Inference in a topically organized semantic net," In *Proc. AAAI-86*, pages 334-338. Philadelphia, PA., Aug. 11-15, 1986.
- [Dowty, 1979] D. Dowty, *Word Meaning and Montague Grammar*, Reidel, Dordrecht, Holland, 1979.
- [Dowty, 1982] D. Dowty, "Tense, time adverbs and compositional semantic theory," *Linguistics and Philosophy*, 5:23-55, 1982.
- [Dowty, 1986] D. Dowty, "The effect of aspectual classes on the temporal structure of discourse: semantics or pragmatics?," *Linguistics and Philosophy*, 9(1):37-61, 1986.
- [Dyer, 1983] M. G. Dyer, *In-Depth Understanding*, MIT Press, Cambridge, Mass., 1983.

- [Enç, 1981] M. Enç, *An Analysis of Nouns as Indexicals*, PhD thesis, U. of Wisconsin, Madison, 1981, Also available from Indiana U. Linguistics Club.
- [Fenstad *et al.*, 1987] J. E. Fenstad, P.-K. Halvorsen, T. Langholm, and J. van Benthem, *Situations, Language and Logic*, D. Reidel Publ., Boston, Mass., 1987.
- [Firby and McDermott, 1987] R. J. Firby and D. McDermott, "Representing and solving temporal planning problems," In N. Cercone and G. McCalla, editors, *The Knowledge Frontier: Essays in the Representation of Knowledge*. Springer-Verlag, New York, 1987.
- [Gazdar *et al.*, 1985] G. Gazdar, E. Klein, G. Pullum, and I. Sag, *Generalized Phrase Structure Grammar*, Harvard Univ. Press, Cambridge, Mass., 1985.
- [Heim, 1982] I. Heim, *The Semantics of Definite and Indefinite Noun Phrases*, PhD thesis, U. of Mass., 1982.
- [Hinrichs, 1988] E. W. Hinrichs, "Tense, quantifiers, and contexts," *Computational Linguistics*, 14(2):3-14, 1988.
- [Hirst, 1989] G. Hirst, "Ontological assumptions in knowledge representation," In *Proc. 1st Inter. Conf. on Principles of Knowledge Representation and Reasoning (KR-89)*, pages 157-169. Toronto, May 15-18, 1989.
- [Hobbs *et al.*, 1986] J. R. Hobbs, W. Croft, T. Davies, D. Edwards, and K. Laws, "Commonsense metaphysics and lexical semantics," In *Proc. of the 24th Annual Meeting of the ACL*, pages 231-240. New York, June 10-13, 1986.
- [Hurum, 1987] S. Hurum, "*Quantifier scoping in initial logical translations of English sentences*," Master's thesis, U. of Alberta, Edmonton, Alberta, 1987.
- [Hurum, 1988] S. Hurum, "Handling scope ambiguities in English," In *Proc. 2nd Conf. on Applied Natural Language Processing*, pages 58-65. Austin, Texas, Feb. 9-12, 1988.
- [Hurum and Schubert, 1986] S. Hurum and L. K. Schubert, "Two types of quantifier scoping," In *Proc. 6th Canadian Conf. on Artificial Intelligence (AI-86)*, pages 39-43. Montreal, May 21-23, 1986.
- [Kamp, 1981] H. Kamp, "A theory of truth and semantic representation," In J. Groenendijk, T. Janssen, and M. Stokhof, editors, *Formal Methods in the Study of Language*. Mathematical Centre-tracts, U. Amsterdam, Amsterdam, 1981.
- [Kamp and Rohrer, 1983] H. Kamp and C. Rohrer, "Tense in Texts," In R. Bauerle, C. Schwarze, and A. von Stechow, editors, *Meaning, Use and Interpretation of Language*, pages 250-269. de Gruyter, Berlin, 1983.

- [Kolodner, 1981] J. L. Kolodner, "Organization and retrieval in a conceptual memory for events or CON54, where are you?," In *Proc. IJCAI-81*, volume 7, pages 227-233. Vancouver, August 24-28, 1981.
- [Kowalski, 1986] R. Kowalski, "Database updates in the Event Calculus," Technical Report DOC 86/12, Dept. of Computing, Imperial College, London, 1986.
- [Kratzer, 1987] A. Kratzer, "An Investigation of the lumps of thought," In J. Pelletier, editor, *Directed Studies on Generics, LI 230*, pages 317-374. Linguistic Institute, Stanford U., June 29 - Aug 7, 1987.
- [Kyburg, 1983] H. E. Kyburg, "The Reference class," *Philosophy of Science*, 50:374-97, 1983.
- [Lehnert *et al.*, 1983] W. G. Lehnert, M. Dyer, P. Johnson, C. Yang, and S. Harley, "BORIS - An experiment in in-depth understanding of narratives," *Artificial Intelligence*, 20:15-62, 1983.
- [McCarthy and Hayes, 1969] J. McCarthy and P. J. Hayes, "Some philosophical problems from the standpoint of artificial intelligence," In B. Meltzer *et al.*, editor, *Machine Intelligence*, 4, pages 463-502. Edinburgh U. Press., 1969.
- [Miller *et al.*, 1987] S. Miller, J. de Haan, and L. K. Schubert, *The User's Guide to ECoNet* (Prepared for Boeing Co. under Purchase Contract W-278258), Edmonton, Alberta, 1987.
- [Miller and Schubert, 1988] S. Miller and L. K. Schubert, "Using specialists to accelerate general reasoning," In *Proc. AAAI-88*, pages 161-165, St. Paul, Minn., August 21-26, 1988.
- [Moens and Steedman, 1988] M. Moens and M. Steedman, "Temporal ontology and temporal reference," *Computational Linguistics*, 14(2):15-28, 1988.
- [Pearl, 1988] J. Pearl, *Probabilistic Reasoning in Intelligent Systems*, Morgan Kaufman, San Mateo, CA, 1988.
- [Pollard and Sag., 1987] C. Pollard and I. A. Sag., *Information-based Syntax and Semantics: Vol. 1. Fundamentals*, CSLI: Center for the Study of Language and Information, Stanford, CA, 1987.
- [Reichenbach, 1947] H. Reichenbach, *Elements of Symbolic Logic*, Macmillan, New York, NY, 1947.
- [Richards and Heny, 1982] B. Richards and F. Heny, "Tense, aspect, and time adverbials, Part I," *Linguistics and Philosophy*, 5:59-107, 1982.

- [Schubert and Hwang, 1989] L. K. Schubert and C. H. Hwang, "An Episodic knowledge representation for Narrative Texts," In *1st Inter. Conf. on Principles of Knowledge Representation and Reasoning (KR89)*, pages 444-458, Toronto, Canada, May 15-18, 1989.
- [Schubert *et al.*, 1989] L. K. Schubert, S. Miller, and C. H. Hwang, *The User's Guide to ECOLOGIC* (Prepared for Boeing Co. under Purchase Contract W-278258), Edmonton, Alberta, 1989.
- [Schubert and Pelletier, 1982] L. K. Schubert and F. J. Pelletier, "From English to logic: context free computation of 'conventional' logical translations," *American J. of Comp. Ling.*, 8:26-44, 1982, Also in *Readings in Natural Language Processing*, B. Grosz, K. Jones and B. Webber, eds., 293-311, Morgan Kaufman, Los Altos, Calif., 1986.
- [Schubert and Pelletier, 1987] L. K. Schubert and F. J. Pelletier, "Problems in the interpretation of the logical form of generics, bare plurals, and mass terms," In E. LePore, editor, *New Directions in Semantics*, pages 387-453. Academic Press, London, 1987.
- [Schubert and Pelletier, 1989] L. K. Schubert and F. J. Pelletier, "Generically speaking, or, using discourse representation theory to interpret generics," In G. Chierchia, B. Partee, and R. Turner, editors, *Property Theory, Type Theory, and Semantics, V.2: Semantic Issues*, pages 193-268. Kluwer Academic Publ., Boston, 1989.
- [ter Meulen, 1988] A. ter Meulen, "Structuring domains for events," 1988, MS, Dept. of Linguistics, U. of Washington, Seattle, WA.
- [Vendler, 1967] Z. Vendler, "Causal Relations," *J. of Philosophy*, 64:704-713, 1967.